Optimization Modulo Theories with Linear Rational Costs

ROBERTO SEBASTIANI and SILVIA TOMASI, DISI, University of Trento, Italy

arXiv:1410.6039v1 [cs.LO] 22 Oct 2014

In the contexts of automated reasoning (AR) and formal verification (FV), important *decision* problems are effectively encoded into Satisfiability Modulo Theories (SMT). In the last decade efficient SMT solvers have been developed for several theories of practical interest (e.g., linear arithmetic, arrays, bit-vectors). Surprisingly, little work has been done to extend SMT to deal with *optimization* problems; in particular, we are not aware of any previous work on SMT solvers able to produce solutions which minimize cost functions over *arithmetical* variables. This is unfortunate, since some problems of interest require this functionality.

In the work described in this paper we start filling this gap. We present and discuss two general procedures for leveraging SMT to handle the minimization of linear rational cost functions, combining SMT with standard minimization techniques. We have implemented the procedures within the MathSAT SMT solver. Due to the absence of competitors in the AR, FV and SMT domains, we have experimentally evaluated our implementation against state-of-the-art tools for the domain of *linear generalized disjunctive programming (LGDP)*, which is closest in spirit to our domain, on sets of problems which have been previously proposed as benchmarks for the latter tools. The results show that our tool is very competitive with, and often outperforms, these tools on these problems, clearly demonstrating the potential of the approach.

Categories and Subject Descriptors: F.4.1 [**Mathematical Logic and Formal Languages**]: Mathematical Logic—*Mechanical theorem proving*

General Terms: Theory, Algorithms, Performance

Additional Key Words and Phrases: Satisfiability Modulo Theories, Automated Reasoning, Optimization

1. INTRODUCTION

In the contexts of automated reasoning (AR) and formal verification (FV), important *decision* problems are effectively encoded into and solved as Satisfiability Modulo Theories (SMT) problems. In the last decade efficient SMT solvers have been developed, that combine the power of modern conflict-driven clause-learning (CDCL) SAT solvers with dedicated decision procedures ($\mathcal{T}$-Solvers) for several first-order theories of practical interest like, e.g., those of equality with uninterpreted functions ($\mathcal{EUF}$), of linear arithmetic over the rationals ($\mathcal{LA}(\mathbb{Q})$) or the integers ($\mathcal{LA}(\mathbb{Z})$), of arrays ($\mathcal{AR}$), of bit-vectors ($\mathcal{BV}$), and their combinations. We refer the reader to [Sebastiani 2007; Barrett et al. 2009] for an overview.

Many SMT-encodable problems of interest, however, may require also the capability of finding models that are *optimal* wrt. some cost function over continuous arithmetical variables. For example, in (SMT-based) *planning with resources* [Wolfman and Weld 1999] a plan for achieving a certain goal must be found which not only fulfills some resource constraints (e.g. on time, gasoline consumption, among others) but that also minimizes the usage of some of such resources; in SMT-based *model checking with timed or hybrid systems* (e.g. [Audemard et al. 2002; Audemard et al. 2005]) you may want to find executions which minimize some parameter (e.g.

Authors' addresses: Roberto Sebastiani (roberto.sebastiani@unitn.it), Silvia Tomasi (silvia.tomasi@disi.unitn.it), DISI, Università di Trento, via Sommarive 9, I-38123 Povo, Trento, Italy. Roberto Sebastiani is supported by Semiconductor Research Corporation (SRC) under GRC Research Project 2012-TJ-2266 WOLF.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© YYYY ACM 1529-3785/YYYY/01-ARTA $15.00
DOI: http://dx.doi.org/10.1145/0000000.0000000

elapsed time), or which minimize/maximize the value of some constant parameter (e.g., a clock timeout value) while fulfilling/violating some property (e.g., minimize the closure time interval of a rail-crossing while preserving safety). This also involves, as particular subcases, problems which are traditionally addressed as *linear disjunctive programming (LDP)* [Balas 1998] or *linear generalized disjunctive programming (LGDP)* [Raman and Grossmann 1994; Sawaya and Grossmann 2012], or as SAT/SMT with Pseudo-Boolean (PB) constraints and (weighted partial) MaxSAT/SMT problems [Roussel and Manquinho 2009; Li and Manyà 2009; Nieuwenhuis and Oliveras 2006; Cimatti et al. 2010; Cimatti et al. 2013a]. Notice that the two latter problems can be encoded into each other.

Surprisingly, little work has been done so far to extend SMT to deal with *optimization* problems [Nieuwenhuis and Oliveras 2006; Cimatti et al. 2010; Sebastiani and Tomasi 2012; Dillig et al. 2012; Cimatti et al. 2013a; Manolios and Papavasileiou 2013] (see §6). In particular, to the best of our knowledge, most such works aim at minimizing cost functions over *Boolean* variables (i.e., SMT with PB cost functions or MaxSMT), whilst we are not aware of any previous work on SMT solvers able to produce solutions which minimize cost functions over *arithmetical* variables. Notice that the former can be encoded into the latter, but not vice versa.

In this this work we start filling this gap. We present two general procedures for adding to SMT the functionality of finding models which minimize some $\mathcal{LA}(\mathbb{Q})$ cost variable —$\mathcal{T}$ being some possibly-empty stably-infinite theory s.t. $\mathcal{T}$ and $\mathcal{LA}(\mathbb{Q})$ are signature disjoint. These two procedures combine standard SMT and minimization techniques: the first, called *offline*, is much simpler to implement, since it uses an incremental SMT solver as a black-box, whilst the second, called *inline*, is more sophisticate and efficient, but it requires modifying the code of the SMT solver. This distinction is important, since the source code of many SMT solvers is not publicly available.

We have implemented these procedures within the MATHSAT5 SMT solver [Cimatti et al. 2013b]. Due to the absence of competitors from AR, FV and SMT domains (§6), we have experimentally evaluated our implementation against state-of-the-art tools for the domain of LGDP, which is closest in spirit to our domain, on sets of problems which have been previously proposed as benchmarks for the latter tools, and on other problem sets. (Notice that LGDP is limited to plain $\mathcal{LA}(\mathbb{Q})$, so that, e.g., it cannot handle combinations of theories like $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$.) The results show that our tool is very competitive with, and often outperforms, these tools on these problems, clearly demonstrating the potential of the approach.

Content. The rest of the paper is organized as follows: in §2 we provide some background knowledge about SAT, SMT, and LGDP; in §3 we formally define the problem addressed, provide the necessary formal results for its solution, and show how the problem generalizes many known optimization problems; in §4 we present our novel procedures; in §5 we present an extensive experimental evaluation; in §6 we survey the related work; in §7 we briefly conclude and highlight directions for future work. In Appendix A we provide the proofs of all the theorems presented in the paper.

Disclaimer. This work was presented in a preliminary form in a much shorter paper at IJCAR 2012 conference [Sebastiani and Tomasi 2012]. Here the content is extended in many ways: first, we provide the theoretical foundations of the procedures, including formal definitions, theorems and relative proofs; second, we provide a much more detailed description and analysis of the procedures, describing in details issues which were only hinted in the conference paper; third, we introduce novel improvements to the procedures; fourth, we provide a much more extended empirical evaluation; finally, we provide a detailed description of the background and of the related work.

2. BACKGROUND

In this section we provide the necessary background about SAT (§2.1), SMT (§2.2), and LGDP (§2.3). We assume a basic background knowledge about logic and operational research. We provide a uniform notation for SAT and SMT: we use boldface lowcase letters $\mathbf{a}, \mathbf{y}$ for arrays and boldface upcase letters $\mathbf{A}, \mathbf{Y}$ for matrices (i.e., two-dimensional arrays), standard lowcase letters

a, y for single rational variables/constants or indices and standard upcase letters A, Y for Boolean atoms and index sets; we use the first five letters in the various forms $\mathbf{a}, ...\mathbf{e}$, ... $A, ...E$, to denote *constant* values, the last five $\mathbf{v}, ...\mathbf{z}$, ... $V, ...Z$ to denote *variables*, and the letters i, j, k, I, J, K for indexes and index sets respectively, subscripts $._j$ denote the j-th element of an array or matrix, whilst superscripts $.^{ij}$ are just indexes, being part of the name of the element. We use lowcase Greek letters $\varphi, \phi, \psi, \mu, \eta$ for denoting formulas and upcase ones Φ, Ψ for denoting sets of formulas.

Remark 2.1. Although we refer to quantifier-free formulas, as it is standard practice in SAT, SMT, CSP and OR communities, with a little abuse of terminology we call "Boolean variables" the propositional atoms and we call "variables" the free constants x_i in quantifier-free $\mathcal{LA}(\mathbb{Q})$-atoms like "$(3x_1 - 2x_2 + x_3 \leq 3)$".

We assume the standard syntactic and semantic notions of propositional logic. Given a non-empty set of primitive propositions $\mathcal{P} = \{p_1, p_2, \ldots\}$, the language of propositional logic is the least set of formulas containing $\mathcal{P}$ and the primitive constants $\top$ and $\bot$ ("true" and "false") and closed under the set of standard propositional connectives $\{\neg, \wedge, \vee, \rightarrow, \leftrightarrow\}$. We call a *propositional atom* every primitive proposition in $\mathcal{P}$, and a *propositional literal* every propositional atom (*positive literal*) or its negation (*negative literal*). We implicitly remove double negations: e.g., if l is the negative literal $\neg p_i$, then by $\neg l$ we mean p_i rather than $\neg\neg p_i$. With a little abuse of notation, we represent a truth assignment μ indifferently either as a *set* of literals $\{l_i\}_i$, with the intended meaning that a positive [resp. negative] literal p_i means that p_i is assigned to true [resp. false], or as a *conjunction* of literals $\bigwedge_i l_i$; thus, e.g., we may say "$l_i \in \mu$" or "$\mu_1 \subseteq \mu_2$", but also "$\neg\mu$" meaning the clause "$\bigvee_i \neg l_i$".

A propositional formula is in *conjunctive normal form, CNF,* if it is written as a conjunction of disjunctions of literals: $\bigwedge_i \bigvee_j l_{ij}$. Each disjunction of literals $\bigvee_j l_{ij}$ is called a *clause*. A *unit clause* is a clause with only one literal.

The above notation and terminology about [positive/negative] literals, truth assignments, CNF and [unit] clauses extend straightforwardly to quantifier-free first-order formulas.

2.1. SAT and CDCL SAT solvers

We present here a brief description on how a Conflict-Driven Clause-Learning (CDCL) SAT solver works. We refer the reader, e.g., to [Marques-Silva and Sakallah 1996; Moskewicz et al. 2001; Marques-Silva et al. 2009] for a detailed description.

We assume the input propositional formula φ is in CNF. (If not, it is first CNF-ized as in [Plaisted and Greenbaum 1986].) The assignment μ is initially empty, and it is updated in a stack-based manner. The SAT solver performs an external loop, alternating three main phases: *Decision*, *Boolean Constraint Propagation (BCP)* and *Backjumping and Learning*.

During *Decision* an unassigned literal l from φ is selected according to some heuristic criterion, and it is pushed into μ. l is called *decision literal* and the number of decision literals which are contained in μ immediately after deciding l is called the *decision level* of l.

Then *BCP* iteratively deduces the literals $l_1, l_2, ...$ deriving from the current assignment and pushes them into μ. *BCP* is based on the iterative application of *unit propagation*: if all but one literals in a clause are false, then the only unassigned literal l is added to μ, all negative occurrences of l in other clauses are declared false and all clauses with positive occurrences of l are declared satisfied. Current SAT solvers include rocket-fast implementations of *BCP* based on the *two-watched-literal scheme*, see [Moskewicz et al. 2001; Marques-Silva et al. 2009]. *BCP* is repeated until either no more literals can be deduced, so that the loop goes back to another decision step, or no more Boolean variable can be assigned, so that the SAT solver ends returning SAT, or μ falsifies some clause ψ of φ (*conflicting clause*).

In the latter case, *Backjumping and Learning* are performed. A process of *conflict analysis* [1] detects a subset η of μ which actually caused the falsification of ψ (*conflict set*) [2] and the decision level blevel where to backtrack. Additionally, the *conflict clause* $\psi' \stackrel{\text{def}}{=} \neg\eta$ is added to φ (*Learning*) and the procedure backtracks up to blevel (*Backjumping*), popping out of μ all literals whose decision level is greater than blevel. When two contradictory literals $l, \neg l$ are assigned at level 0, the loop terminates, returning UNSAT.

Notice that CDCL SAT solvers implement "safe" strategies for deleting clauses when no more necessary, which guarantee the use of polynomial space without affecting the termination, correctness and completeness of the procedure. (See e.g. [Marques-Silva et al. 2009; Nieuwenhuis et al. 2006].)

Many modern CDCL SAT solvers provide a *stack-based incremental interface* (see e.g. [Eén and Sörensson 2004]), by which it is possible to push/pop sub-formulas ϕ_i into a stack of formulas $\Phi \stackrel{\text{def}}{=} \{\phi_1, ..., \phi_k\}$, and check incrementally the satisfiability of $\bigwedge_{i=1}^{k} \phi_i$. The interface maintains most of the information about the *status* of the search from one call to the other, in particular it records the learned clauses (plus other information). Consequently, when invoked on Φ the solver can reuse a clause ψ which was learned during a previous call on some Φ' if ψ was derived only from clauses which are still in Φ —provided ψ was not discharged in the meantime; in particular, if $\Phi' \subseteq \Phi$, then the solver can reuse all clauses learned while solving Φ'.

Another important feature of many incremental CDCL SAT solvers is their capability, when Φ is found unsatisfiable, to return a subset of formulas in Φ which caused the unsatisfiability of Φ. This is related to the problem of finding an *unsatisfiable core* of a formula, see e.g. [Lynce and Marques-Silva 2004]. Notice that such subset is not unique, and it is not necessarily minimal.

2.2. SMT and Lazy SMT solvers

We assume a basic background knowledge on first-order logic. to *ground* formulas/literals/atoms in the language of $\mathcal{T}$ ($\mathcal{T}$-formulas/literals/atoms hereafter). Notice that, for better readability, with a little abuse of notation we often refer to a theory $\mathcal{T}$ instead of its corresponding signature; also by "empty theory" we mean the empty theory over the empty signature; finally, by adopting the terminology in Remark 2.1, we say that a *variable* belongs to the signature of a theory $\mathcal{T}$ (or simply that it belongs to a theory $\mathcal{T}$).

A *theory solver for* $\mathcal{T}$, $\mathcal{T}$-Solver, is a procedure able to decide the $\mathcal{T}$-satisfiability of a conjunction/set μ of $\mathcal{T}$-literals. If μ is $\mathcal{T}$-unsatisfiable, then $\mathcal{T}$-Solver returns UNSAT and a set/conjunction η of $\mathcal{T}$-literals in μ which was found $\mathcal{T}$-unsatisfiable; η is called a *$\mathcal{T}$-conflict set*, and $\neg\eta$ a *$\mathcal{T}$-conflict clause*. If μ is $\mathcal{T}$-satisfiable, then $\mathcal{T}$-Solver returns SAT; it may also be able to return some unassigned $\mathcal{T}$-literal $l \notin \mu$ from a set of all available $\mathcal{T}$-literals, s.t. $\{l_1, ..., l_n\} \models_{\mathcal{T}} l$, where $\{l_1, ..., l_n\} \subseteq \mu$. We call this process *$\mathcal{T}$-deduction* and $(\bigvee_{i=1}^{n} \neg l_i \vee l)$ a *$\mathcal{T}$-deduction clause*. Notice that $\mathcal{T}$-conflict and $\mathcal{T}$-deduction clauses are valid in $\mathcal{T}$. We call them *$\mathcal{T}$-lemmas*.

Given a $\mathcal{T}$-formula φ, the formula φ^p obtained by rewriting each $\mathcal{T}$-atom in φ into a fresh atomic proposition is the *Boolean abstraction* of φ, and φ is the *refinement* of φ^p. Notationally, we indicate by φ^p and μ^p the Boolean abstraction of φ and μ, and by φ and μ the refinements of φ^p and μ^p respectively. With a little abuse of notation, we say that μ^p is $\mathcal{T}$-(un)satisfiable iff μ is $\mathcal{T}$-(un)satisfiable. We say that the truth assignment μ *propositionally satisfies* the formula φ, written $\mu \models_p \varphi$, if $\mu^p \models \varphi^p$.

In a lazy SMT($\mathcal{T}$) solver, the Boolean abstraction φ^p of the input formula φ is given as input to a CDCL SAT solver, and whenever a satisfying assignment μ^p is found s.t. $\mu^p \models \varphi^p$, the corresponding set of $\mathcal{T}$-literals μ is fed to the $\mathcal{T}$-Solver; if μ is found $\mathcal{T}$-consistent, then φ is $\mathcal{T}$-consistent;

[1] When a clause ψ is falsified by the current assignment a *conflict clause* ψ' is computed from ψ s.t. ψ' contains only one literal l_u which has been assigned at the last decision level. ψ' is computed starting from $\psi' = \psi$ by iteratively resolving ψ' with the clause ψ_l causing the unit-propagation of some literal l in ψ' until some stop criterion is met.

[2] That is, η is enough to force the unit-propagation of the literals causing the failure of ψ.

otherwise, $\mathcal{T}$-Solver returns a $\mathcal{T}$-conflict set η causing the inconsistency, so that the clause $\neg\eta^p$ is used to drive the backjumping and learning mechanism of the SAT solver. The process proceeds until either a $\mathcal{T}$-consistent assignment μ is found, or no more assignments are available (φ is $\mathcal{T}$-inconsistent).

Important optimizations are *early pruning* and *$\mathcal{T}$-propagation*. The $\mathcal{T}$-Solver is invoked also when an assignment μ is still under construction: if it is $\mathcal{T}$-unsatisfiable, then the procedure backtracks, without exploring the (possibly many) extensions of μ; if not, and if the $\mathcal{T}$-Solver is able to perform a $\mathcal{T}$-deduction $\{l_1, ..., l_n\} \models_{\mathcal{T}} l$, then l can be unit-propagated, and the $\mathcal{T}$-deduction clause $(\bigvee_{i=1}^{n} \neg l_i \vee l)$ can be used in backjumping and learning. To this extent, in order to maximize the efficiency, most $\mathcal{T}$-solvers are *incremental* and *backtrackable*, that is, they are called via a push&pop interface, maintaining and reusing the status of the search from one call and the other.

Another optimization is *pure-literal filtering*: if some $\mathcal{LA}(\mathbb{Q})$-atoms occur only positively [resp. negatively] in the original formula (learned clauses are ignored), then we can safely drop every negative [resp. positive] occurrence of them from the assignment μ to be checked by the $\mathcal{T}$-Solver [Sebastiani 2007]. Intuitively, since such occurrences play no role in satisfying the formula, the resulting partial assignment $\mu^{p\prime}$ still satisfies φ^p. The benefits of this action are twofold: (i) it reduces the workload for the $\mathcal{T}$-Solver by feeding to it smaller sets; (ii) it increases the chance of finding a $\mathcal{T}$-consistent satisfying assignment by removing "useless" $\mathcal{T}$-literals which may cause the $\mathcal{T}$-inconsistency of μ.

The above schema is a coarse abstraction of the procedures underlying all the state-of-the-art lazy SMT tools. The interested reader is pointed to, e.g., [Nieuwenhuis et al. 2006; Sebastiani 2007; Barrett et al. 2009] for details and further references. Importantly, some SMT solvers, including MATHSAT, inherit from their embedded SAT solver the capabilities of working incrementally and of returning the subset of input formulas causing the inconsistency, as described in §2.1.

The *Theory of Linear Arithmetic on the rationals* ($\mathcal{LA}(\mathbb{Q})$) and on the integers ($\mathcal{LA}(\mathbb{Z})$) is one of the theories of main interest in SMT. It is a first-order theory whose atoms are of the form $(a_1x_1 + \ldots + a_nx_n \diamond b)$, i.e. $(\mathbf{ax} \diamond b)$, s.t $\diamond \in \{=, \neq, <, >, \leq, \geq\}$.

Efficient incremental and backtrackable procedures have been conceived in order to decide $\mathcal{LA}(\mathbb{Q})$ [Dutertre and de Moura 2006] and $\mathcal{LA}(\mathbb{Z})$ [Griggio 2012]. In particular, for $\mathcal{LA}(\mathbb{Q})$ most SMT solvers implement variants of the simplex-based algorithm by Dutertre and de Moura [Dutertre and de Moura 2006] which is specifically designed for integration in a lazy SMT solver, since it is fully incremental and backtrackable and allows for aggressive $\mathcal{T}$-deduction. Another benefit of such algorithm is that it handles *strict inequalities* directly. Its method is based on the fact that a set of $\mathcal{LA}(\mathbb{Q})$ atoms Γ containing strict inequalities $S = \{0 < t_1, \ldots, 0 < t_n\}$ is satisfiable iff there exists a rational number $\epsilon > 0$ such that $\Gamma_\epsilon \stackrel{\text{def}}{=} (\Gamma \cup S_\epsilon) \setminus S$ is satisfiable, s.t. $S_\epsilon \stackrel{\text{def}}{=} \{\epsilon \leq t_1, \ldots, \epsilon \leq t_n\}$. The idea of [Dutertre and de Moura 2006] is that of treating the *infinitesimal parameter* ϵ symbolically instead of explicitly computing its value. Strict bounds $(x < b)$ are replaced with weak ones $(x \leq b - \epsilon)$, and the operations on bounds are adjusted to take ϵ into account. We refer the reader to [Dutertre and de Moura 2006] for details.

2.3. Linear Generalized Disjunctive Programming

Mixed Integer Linear Programming (MILP) is an extension of Linear Programming (LP) involving both discrete and continuous variables [Lodi 2009]. MILP problems have the following form:

$$\min\{\mathbf{cx} : \mathbf{Ax} \geq \mathbf{b}, \mathbf{x} \geq 0, \mathbf{x}_j \in \mathbb{Z}\ \forall j \in I\} \tag{1}$$

where $\mathbf{A}$ is a matrix, $\mathbf{c}$ and $\mathbf{b}$ are constant vectors and $\mathbf{x}$ the variable vector. A large variety of techniques and tools for MILP are available, mostly based on efficient combinations of LP, *branch-and-bound* search mechanism and *cutting-plane* methods, resulting in a *branch-and-cut* approach (see e.g. [Lodi 2009]). SAT techniques have also been incorporated into these procedures for MILP (see [Achterberg et al. 2008]).

The branch-and-bound search iteratively partitions the solution space of the original MILP problem into subproblems and solves their LP relaxation (i.e. a MILP problem where the integrality constraint on the variables $\mathbf{x}_j$, for all $j \in I$, is dropped) until all variables are integral in the optimal solution of the LP relaxation. Cutting planes (e.g. Gomory mixed-integer and lift-and-project cuts [Lodi 2009]) are linear inequalities that can be inferred and added to the original MILP problem and its subproblems in order to cut away non-integer solutions of the LP relaxation and obtain tighter relaxations.

Linear Disjunctive Programming (LDP) problems are LP problems where linear constraints are connected by the logical operations of conjunction and disjunction (see, e.g., [Balas 1998]). The constraint set can be expressed by a disjunction of linear systems (*Disjunctive Normal Form*):

$$\bigvee_{i \in I} (\mathbf{A}^i \mathbf{x} \geq \mathbf{b}^i) \tag{2}$$

or, alternatively, as a conjunction (*Conjunctive Normal Form*):

$$(\mathbf{A}\mathbf{x} \geq \mathbf{b}) \wedge \bigwedge_{j=1}^{t} \bigvee_{k \in I_j} (\mathbf{c}^k \mathbf{x} \geq d^k) \tag{3}$$

or in an intermediate form called *Regular Form* (see, e.g., [Balas 1983]). Notice that (3) can be obtained from (2) by factoring out the common inequalities $(\mathbf{A}\mathbf{x} \geq \mathbf{b})$ and then by applying the distributivity of $\wedge$ and $\vee$, although the latter step can cause a blowup in size. LDP problems are effectively solved by the lift-and-project approach which combines a family of cutting planes, called lift-and-project cuts, and the branch-and-bound schema (see, e.g., [Balas and Bonami 2007]).

Linear Generalized Disjunctive Programming (LGDP), is a generalization of LDP which has been proposed in [Raman and Grossmann 1994] as an alternative model to the MILP problem. Unlike MILP, which is based entirely on algebraic equations and inequalities, the LGDP model allows for combining algebraic and logical equations with Boolean propositions through Boolean operations, providing a much more natural representation of discrete decisions. Current approaches successfully address LGDP by reformulating and solving it as a MILP problem [Raman and Grossmann 1994; Vecchietti and Grossmann 2004; Sawaya and Grossmann 2005; Sawaya and Grossmann 2012]; these reformulations focus on efficiently encoding disjunctions and logic propositions into MILP, so as to be fed to an efficient MILP solver like CPLEX.

The general formulation of a LGDP problem is the following [Raman and Grossmann 1994]:

$$\begin{array}{ll} \min & \sum_{k \in K} \mathbf{z}_k + \mathbf{d}\mathbf{x} \\ \text{s.t.} & \mathbf{B}\mathbf{x} \leq \mathbf{b} \\ & \bigvee_{j \in J_k} \begin{bmatrix} Y^{jk} \\ \mathbf{A}^{jk}\mathbf{x} \leq \mathbf{a}^{jk} \\ \mathbf{z}_k = c^{jk} \end{bmatrix} \qquad \forall k \in K \\ & \phi \\ & \mathbf{0} \leq \mathbf{x} \leq \mathbf{e} \\ & \mathbf{z}_k, c^{jk} \in \mathbb{R}^1_+, Y^{jk} \in \{True, False\} \ \forall j \in J_k, \forall k \in K \end{array} \tag{4}$$

where $\mathbf{x}$ is a vector of positive rational variables, $\mathbf{d}$ is a vector of positive rational values representing the cost-per-unit of each variable in $\mathbf{x}$, $\mathbf{z}$ is a vector of positive rational variables representing the cost assigned to each disjunction, c^{jk} are positive constant values, $\mathbf{e}$ is a vector of upper bounds for $\mathbf{x}$ and Y^{jk} are Boolean variables.

The disequalities $\mathbf{B}\mathbf{x} \leq \mathbf{b}$, where $(\mathbf{B}, \mathbf{b})$ is a $m \times (n+1)$ matrix, are the "common" constraints that must always hold.

Each disjunction $k \in K$ consists in at least two disjuncts $j \in J_k$, s.t. the jk-th disjunct contains:

(i) the Boolean variable Y^{jk}, representing discrete decisions,
(ii) a set of linear constraints $\mathbf{A}^{jk}\mathbf{x} \leq \mathbf{a}^{jk}$, where $(\mathbf{A}^{jk}, \mathbf{a}^{jk})$ is a $m_{jk} \times (n+1)$ matrix,
(iii) the equality $\mathbf{z}_k = c^{jk}$, assigning the value c^{jk} to the cost variable $\mathbf{z}_k$.

Each disjunct is true if and only if all the three elements (i)-(iii) above are true. ϕ is a propositional formula, expressed in Conjunctive Normal Form, which must contain the "xor" constraints $\bigoplus_{j \in J_k} Y^{jk}$ for each $k \in K$, plus possibly other constraints. Intuitively, for each $k \in K$, the only variable Y^{jk} which is set to true selects the set of disequalities $\mathbf{A}^{jk}\mathbf{x} \leq \mathbf{a}^{jk}$ which are enforced and hence it selects the relative cost c^{jk} of this choice to be assigned to the cost variable $\mathbf{z}_k$.

LGDP problems can be solved using MILP solvers by reformulating the original problem in different ways; big-M (BM) and convex hull (CH) are the two most common reformulations. In BM, the Boolean variables Y^{jk} and the logic constraints ϕ are replaced by binary variables $\mathbf{Y}_{jk}$ and linear inequalities as follows [Raman and Grossmann 1994]:

$$\begin{array}{lcll} \min & \sum_{k \in K} \sum_{j \in J_k} c^{jk} \mathbf{Y}_{jk} + \mathbf{dx} & & \\ \text{s.t.} & \mathbf{Bx} \leq \mathbf{b} & & \\ & \mathbf{A}^{jk}\mathbf{x} - \mathbf{a}^{jk} \leq \mathbf{M}^{jk}(1 - \mathbf{Y}_{jk}) & \forall j \in J_k, \forall k \in K & \\ & \sum_{j \in J_k} \mathbf{Y}_{jk} = 1 & \forall k \in K & (5) \\ & \mathbf{DY} \leq \mathbf{D}' & & \\ & \mathbf{x} \in \mathbb{R}^n \; s.t. \; \mathbf{0} \leq \mathbf{x} \leq \mathbf{e}, \mathbf{Y}_{jk} \in \{0,1\} & \forall j \in J_k, \forall k \in K & \end{array}$$

where $\mathbf{M}^{jk}$ are the '"big-M" parameters that makes redundant the system of constraint $j \in J_k$ in the disjunction $k \in K$ when $\mathbf{Y}_{jk} = 0$ and the constraints $\mathbf{DY} \leq \mathbf{D}'$ are derived from ϕ.

In CH, the Boolean variables Y^{jk} are replaced by binary variables $\mathbf{Y}_{jk}$ and the variables $\mathbf{x} \in \mathbb{R}^n$ are disaggregated into new variables $\mathbf{v} \in \mathbb{R}^n$ in the following way:

$$\begin{array}{lcll} \min & \sum_{k \in K} \sum_{j \in J_k} c^{jk} \mathbf{Y}_{jk} + \mathbf{dx} & & \\ \text{s.t.} & \mathbf{Bx} \leq \mathbf{b} & & \\ & \mathbf{A}^{kj}\mathbf{v}^{jk} \leq \mathbf{a}^{jk}\mathbf{Y}_{jk} & \forall j \in J_k, \forall k \in K & \\ & \mathbf{x} = \sum_{j \in J_k} \mathbf{v}_{jk} & \forall k \in K & (6) \\ & \mathbf{v}_{jk} \leq \mathbf{Y}_{jk}\mathbf{e}^{jk} & \forall j \in J_k, \forall k \in K & \\ & \sum_{j \in J_k} \mathbf{Y}_{jk} = 1 & \forall k \in K & \\ & \mathbf{DY} \leq \mathbf{D}' & & \\ & \mathbf{x}, \mathbf{v} \in \mathbb{R}^n \; s.t. \; \mathbf{0} \leq \mathbf{x}, \mathbf{v}, \; \mathbf{Y}_{jk} \in \{0,1\} & \forall j \in J_k, \forall k \in K & \end{array}$$

where constant $\mathbf{e}^{jk}$ are upper bounds for variables $\mathbf{v}$ chosen to match the upper bounds on the variables $\mathbf{x}$.

Sawaya and Grossman [Sawaya and Grossmann 2005] observed two facts. First, the relaxation of BM is often weak causing a higher number of nodes examined in the branch-and-bound search. Second, the disaggregated variables and new constraints increase the size of the reformulation leading to a high computational effort. In order to overcome these issues, they proposed a cutting plane method that consists in solving a sequence of BM relaxations with cutting planes that are obtained from CH relaxations. They provided an evaluation of the presented algorithm on three different problems: strip-packing, retrofit planning and zero-wait job-shop scheduling problems.

3. OPTIMIZATION IN SMT($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$)

In this section we define the problem addressed (§3.1), we introduce the formal foundations for its solution (§3.2), and we show how it generalizes many known optimization problems from the literature (§3.3).

3.1. Basic Definitions and Notation

In this paper we consider only signature-disjoint stably-infinite theories with equality $\mathcal{T}_i$ ("Nelson-Oppen theories" [Nelson and Oppen 1979]) and we focus our interest on $\mathcal{LA}(\mathbb{Q})$. In particular, in what follows we assume $\mathcal{T}$ to be some stably-infinite theory with equality, s.t. $\mathcal{LA}(\mathbb{Q})$ and $\mathcal{T}$ are signature-disjoint. $\mathcal{T}$ can also be a combination of Nelson-Oppen theories.

We assume the standard model of $\mathcal{LA}(\mathbb{Q})$, whose domain is the set of rational numbers $\mathbb{Q}$.

Definition 3.1 (*OMT*($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$), *OMT*($\mathcal{LA}(\mathbb{Q})$), *and* $\mathsf{min}_{\mathsf{cost}}$.). Let φ be a ground SMT($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$) formula, and cost be a $\mathcal{LA}(\mathbb{Q})$ variable occurring in φ. We call an *Optimization Modulo* $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$ *problem*, written *OMT*($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$), the problem of finding a model for φ (if any) whose value of cost is minimum. We denote such value as $\mathsf{min}_{\mathsf{cost}}(\varphi)$. If φ is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-unsatisfiable, then $\mathsf{min}_{\mathsf{cost}}(\varphi)$ is $+\infty$; if there is no minimum value for cost, then $\mathsf{min}_{\mathsf{cost}}(\varphi)$ is $-\infty$. We call an *Optimization Modulo* $\mathcal{LA}(\mathbb{Q})$ problem, written *OMT*($\mathcal{LA}(\mathbb{Q})$), an OMT($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$) problem where $\mathcal{T}$ is the empty theory.

A dual definition where we look for a *maximum* value is easy to formulate.

In order to make the discussion simpler, we assume w.l.o.g. that all $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$ formulas are *pure* [Nelson and Oppen 1979]. With a little abuse of notation, we say that an atom in a ground $\mathcal{T}_1 \cup \mathcal{T}_2$ formula is $\mathcal{T}_i$-pure if it contains only variables and symbols from the signature of $\mathcal{T}_i$, for every $i \in \{1, 2\}$; a $\mathcal{T}_1 \cup \mathcal{T}_2$ ground formula is pure iff all its atoms are either $\mathcal{T}_1$-pure or $\mathcal{T}_2$-pure. Although the purity assumption is not necessary (see [Barrett et al. 2002]), it simplifies the explanation, since it allows us to speak about "$\mathcal{LA}(\mathbb{Q})$-atoms" or "$\mathcal{T}$-atoms" without further specifying. Moreover, every non-pure formula can be easily purified [Nelson and Oppen 1979].

We also assume w.l.o.g. that all $\mathcal{LA}(\mathbb{Q})$-atoms containing the variable cost are in the form $(t \bowtie \mathsf{cost})$, s.t. $\bowtie \in \{=, \leq, \geq, <, >\}$ and cost does not occur in t.

Definition 3.2 (*Bounds and range for* cost). If φ is in the form $\varphi' \wedge (\mathsf{cost} < c)$ [resp. $\varphi' \wedge \neg(\mathsf{cost} < c)$] for some value $c \in \mathbb{Q}$, then we call c an *upper bound* [resp. *lower bound*] for cost. If ub [resp lb] is the minimum upper bound [resp. the maximum lower bound] for φ, we also call the interval $[\mathsf{lb}, \mathsf{ub}[$ the *range* of cost.

We adopt the convention of defining upper bounds to be strict and lower bounds to be non-strict for a practical reason: typically an upper bound $(\mathsf{cost} < c)$ derives from the fact that a model $\mathcal{I}$ of cost c has been previously found, whilst a lower bound $\neg(\mathsf{cost} < c)$ derives either from the user's knowledge (e.g. "the cost cannot be lower than zero") of from the fact that the formula $\varphi \wedge (\mathsf{cost} < c)$ has been previously found $\mathcal{T}$-unsatisfiable whilst φ has not.

3.2. Theoretical Results

We present here the theoretical foundations of our procedures. The proofs of the novel results are reported in Appendix A.

The following facts follow straightforwardly from Definition 3.1.

PROPOSITION 3.3. *Let $\varphi, \varphi_1, \varphi_2$ be $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-formulas and μ_1, μ_2 be truth assignments.*

(a) If $\varphi_1 \models \varphi_2$, then $\mathsf{min}_{\mathsf{cost}}(\varphi_1) \geq \mathsf{min}_{\mathsf{cost}}(\varphi_2)$.
(b) If $\mu_1 \supseteq \mu_2$, then $\mathsf{min}_{\mathsf{cost}}(\mu_1) \geq \mathsf{min}_{\mathsf{cost}}(\mu_2)$.
(c) φ is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable if and only if $\mathsf{min}_{\mathsf{cost}}(\varphi) < +\infty$.

We recall first some definitions and results from the literature.

Definition 3.4. We say that a collection $\mathcal{M} := \{\mu_1, \ldots, \mu_n\}$ of (possibly partial) assignments propositionally satisfying φ is *complete* iff, for every total assignment η s.t. $\eta \models_p \varphi$, there exists $\mu_j \in \mathcal{M}$ s.t. $\mu_j \subseteq \eta$.

THEOREM 3.5 ([SEBASTIANI 2007]). *Let φ be a $\mathcal{T}$-formula and let $\mathcal{M} := \{\mu_1, \ldots, \mu_n\}$ be a complete collection of (possibly partial) truth assignments propositionally satisfying φ. Then, φ is $\mathcal{T}$-satisfiable if and only if μ_j is $\mathcal{T}$-satisfiable for some $\mu_j \in \mathcal{M}$.*

Theorem 3.5 is the theoretical foundation of the lazy SMT approach described in §2.2, where a CDCL SAT solver enumerates a complete collection $\mathcal{M}$ of truth assignments as above, whose $\mathcal{T}$-satisfiability is checked by a $\mathcal{T}$-Solver. Notice that in Theorem 3.5 the theory $\mathcal{T}$ can be any combination of theories $\mathcal{T}_i$, including $\mathcal{LA}(\mathbb{Q})$.

Here we extend Theorem 3.5 to OMT($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$) as follows.

THEOREM 3.6. *Let φ be a $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-formula and let $\mathcal{M} \stackrel{\text{def}}{=} \{\mu_1, \ldots, \mu_n\}$ be a complete collection of (possibly-partial) truth assignments which propositionally satisfy φ. Then* $\mathsf{min}_{\mathsf{cost}}(\varphi) = min_{\mu \in \mathcal{M}} \mathsf{min}_{\mathsf{cost}}(\mu)$.

Notice that we implicitly define $min_{\mu \in \mathcal{M}} \mathsf{min}_{\mathsf{cost}}(\mu) \stackrel{\text{def}}{=} +\infty$ if $\mathcal{M}$ is empty. Since $\mathsf{min}_{\mathsf{cost}}(\mu)$ is $+\infty$ if μ is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-unsatisfiable, we can safely restrict the search for minima to the $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable assignments in $\mathcal{M}$.

If $\mathcal{T}$ is the empty theory, then the notion of $\mathsf{min}_{\mathsf{cost}}(\mu)$ is straightforward, since each μ is a conjunction of Boolean literals and of $\mathcal{LA}(\mathbb{Q})$ constraints, so that Theorem 3.6 provides the theoretical foundation for OMT($\mathcal{LA}(\mathbb{Q})$).

If instead $\mathcal{T}$ is not the empty theory, then each μ is a set of Boolean literals and of pure $\mathcal{T}$-literals and $\mathcal{LA}(\mathbb{Q})$ constraints sharing variables, so that the notion of $\mathsf{min}_{\mathsf{cost}}(\mu)$ is not straightforward. To cope with this fact, we first recall from the literature some definitions and an important result.

Definition 3.7 (*Interface variables, interface equalities.*). Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be two stably-infinite theories with equality and disjoint signatures, and let φ be a $\mathcal{T}_1 \cup \mathcal{T}_2$-formula. We call *interface variables* of φ the variables occurring in both $\mathcal{T}_1$-pure and $\mathcal{T}_2$-pure atoms of φ, and *interface equalities* of φ the equalities $(x_i = x_j)$ on the interface variables of φ.

As it is common practice in SMT (see e.g. [Tinelli and Harandi 1996]) hereafter we consider only interface equalities modulo reflexivity and symmetry, that is, we implicitly assume some total order $\preceq$ on the interface variables x_i of φ, and we restrict w.l.o.g. the set of interface equalities on φ to $\mathcal{IE}(\varphi) \stackrel{\text{def}}{=} \{(x_i = x_j) \mid x_i \prec x_j\}$, dropping thus uninformative equalities like $(x_i = x_i)$ and considering only the first equality in each pair $\{(x_i = x_j), (x_j = x_i)\}$.

Notation-wise, in what follows we use the subscripts e, d, i in "$\mu_{\ldots}$", like in "μ_{ed}", to denote conjunctions of equalities, disequalities and inequalities between interface variables respectively.

THEOREM 3.8 ([TINELLI AND HARANDI 1996]). *Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be two stably-infinite theories with equality and disjoint signatures; let $\mu \stackrel{\text{def}}{=} \mu_{\mathcal{T}_1} \wedge \mu_{\mathcal{T}_2}$ be a conjunction of $\mathcal{T}_1 \cup \mathcal{T}_2$-literals s.t. each $\mu_{\mathcal{T}_i}$ is pure for $\mathcal{T}_i$. Then μ is $\mathcal{T}_1 \cup \mathcal{T}_2$-satisfiable if and only if there exists an equivalence class $Eq \subseteq \mathcal{IE}(\mu)$ over the interface variables of μ and the corresponding total truth assignment μ_{ed} to the interface equalities over μ:*

$$\mu_{ed} \stackrel{\text{def}}{=} \mu_e \wedge \mu_d, \;\; s.t.\; \mu_e \stackrel{\text{def}}{=} \bigwedge_{(x_i, x_j) \in Eq} (x_i = x_j), \;\; \mu_d \stackrel{\text{def}}{=} \bigwedge_{(x_i, x_j) \notin Eq} \neg(x_i = x_j) \tag{7}$$

s.t. $\mu_{\mathcal{T}_k} \wedge \mu_{ed}$ is $\mathcal{T}_k$-satisfiable for every $k \in \{1, 2\}$.

Theorem 3.8 is the theoretical foundation of, among others, the *Delayed Theory Combination* SMT technique for combined theories [Bozzano et al. 2006], where a CDCL SAT solver enumerates a complete collection of extended assignments $\mu \wedge \mu_{ed}$, which propositionally satisfy the input

formula, and dedicated $\mathcal{T}_k$-solvers check independently the $\mathcal{T}_k$-satisfiability of $\mu_{\mathcal{T}_k} \wedge \mu_{ed}$, for each $k \in \{1, 2\}$.

We consider now a $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$ formula φ and a (possibly-partial) truth assignment μ which propositionally satisfies it. μ can be written as $\mu \stackrel{\text{def}}{=} \mu_{\mathbb{B}} \wedge \mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{\mathcal{T}}$, s.t. $\mu_{\mathbb{B}}$ is a consistent conjunction of Boolean literals, $\mu_{\mathcal{LA}(\mathbb{Q})}$ and $\mu_{\mathcal{T}}$ are $\mathcal{LA}(\mathbb{Q})$-pure and $\mathcal{T}$-pure conjunctions of literals respectively. (Notice that the $\mu_{\mathcal{B}}$ component does not affect the $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiability of μ.) Then the following definitions and theorems show how $\mathsf{min}_{\mathsf{cost}}(\mu)$ can be defined and computed.

Definition 3.9. Let $\mu \stackrel{\text{def}}{=} \mu_{\mathbb{B}} \wedge \mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{\mathcal{T}}$ be a truth assignment satisfying some $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$ ground formula, s.t. $\mu_{\mathbb{B}}$ is a consistent conjunction of Boolean literals, $\mu_{\mathcal{LA}(\mathbb{Q})}$ and $\mu_{\mathcal{T}}$ are $\mathcal{LA}(\mathbb{Q})$-pure and $\mathcal{T}$-pure conjunctions of literals respectively. We call the *complete set of* ed*-extensions of* μ the set $\mathcal{EX}_{ed}(\mu) \stackrel{\text{def}}{=} \{\eta_1, ..., \eta_n\}$ of all possible assignments in the form $\mu \wedge \mu_{ed}$, where μ_{ed} is in the form (7), for every equivalence class Eq in $\mathcal{IE}(\mu)$.

THEOREM 3.10. *Let* μ *be as in Definition 3.9. Then*

(a) $\mathsf{min}_{\mathsf{cost}}(\mu) = min_{\eta \in \mathcal{EX}_{ed}(\mu)} \mathsf{min}_{\mathsf{cost}}(\eta)$

(b) forall $\eta \in \mathcal{EX}_{ed}(\mu)$,

$$\mathsf{min}_{\mathsf{cost}}(\eta) = \begin{cases} +\infty & \text{if } \mu_{\mathcal{T}} \wedge \mu_{ed} \text{ is } \mathcal{T}\text{-unsatisfiable or} \\ & \text{if } \mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ed} \text{ is } \mathcal{LA}(\mathbb{Q})\text{-unsatisfiable} \\ \mathsf{min}_{\mathsf{cost}}(\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ed}) & \text{otherwise.} \end{cases}$$

We notice that, at least in principle, computing $\mathsf{min}_{\mathsf{cost}}(\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ed})$ is an operation which can be performed by standard linear-programming techniques (see §4). Thus, by combining Theorems 3.6 and 3.10 we have a general method for computing $\mathsf{min}_{\mathsf{cost}}(\varphi)$ also in the general case of non-empty theory $\mathcal{T}$.

In practice, however, it is often the case that $\mathcal{LA}(\mathbb{Q})$-solvers/optimizers cannot handle efficiently negated equalities like, e.g., $\neg(x_i = x_j)$ (see [Dutertre and de Moura 2006]). Thus, a technique which is adopted by most SMT solver is to expand them into the corresponding disjunction of strict inequalities $(x_i < x_j) \vee (x_i > x_j)$. This "case split" is typically efficiently handled directly by the embedded SAT solver.

We notice, however, that such case-split may be applied also to interface equalities $(x_i = x_j)$, and that the resulting "interface inequalities" $(x_i < x_j)$ and $(x_i > x_j)$ cannot be handled by the other theory $\mathcal{T}$, because "$<$" and "$>$" are $\mathcal{LA}(\mathbb{Q})$-specific symbols. In order to cope with this fact, some more theoretical discussion is needed.

Definition 3.11. Let μ be as in Definition 3.9. We call the *complete set of* edi*-extensions of* μ the set $\mathcal{EX}_{edi}(\mu) \stackrel{\text{def}}{=} \{\rho_1, ..., \rho_n\}$ of all possible truth assignments in the form $\mu \wedge \mu_{ed} \wedge \mu_i$, where μ_{ed} is as in Definition 3.9 and μ_i is a total truth assignment to the atoms in $\{(x_i < x_j), (x_i > x_j) \,|(x_i = x_j) \in \mathcal{IE}(\mu)\}$ s.t. $\mu_{ed} \wedge \mu_i$ is $\mathcal{LA}(\mathbb{Q})$-consistent.

μ_i assigns both $(x_i < x_j)$ and $(x_i > x_j)$ to false if $(x_i = x_j)$ is true in μ_{ed}, one of them to true and the other to false if $(x_i = x_j)$ is false in μ_{ed}. Intuitively, the presence of each negated interface equalities $\neg(x_i = x_j)$ in μ_{ed} forces the choice of one of the two parts $\langle (x_i < x_j), (x_i > x_j) \rangle$ of the solution space.

THEOREM 3.12. *Let* μ *be as in Definition 3.9. Then*

(a) μ *is* $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$*-satisfiable iff some* $\rho \in \mathcal{EX}_{edi}(\mu)$ *is* $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$*-satisfiable.*

(b) $\mathsf{min}_{\mathsf{cost}}(\mu) = min_{\rho \in \mathcal{EX}_{edi}(\mu)} \mathsf{min}_{\mathsf{cost}}(\rho)$.

(c) forall $\rho \in \mathcal{EX}_{edi}(\mu)$, ρ *is* $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$*-satisfiable iff* $\mu_{\mathcal{T}} \wedge \mu_{ed}$ *is* $\mathcal{T}$*-satisfiable and* $\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_e \wedge \mu_i$ *is* $\mathcal{LA}(\mathbb{Q})$*-satisfiable.*

(d) forall $\rho \in \mathcal{EX}_{edi}(\mu)$,

$$\mathsf{min}_{\mathsf{cost}}(\rho) = \begin{cases} +\infty & \textit{if } \mu_{\mathcal{T}} \wedge \mu_{ed} \textit{ is } \mathcal{T}\textit{-unsatisfiable or} \\ & \textit{if } \mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_e \wedge \mu_i \textit{ is } \mathcal{LA}(\mathbb{Q})\textit{-unsatisfiable} \\ \mathsf{min}_{\mathsf{cost}}(\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_e \wedge \mu_i) & \textit{otherwise.} \end{cases}$$

Thus, by combining Theorems 3.6 and 3.12 we have a general method for computing $\mathsf{min}_{\mathsf{cost}}(\varphi)$ in the case of non-empty theory $\mathcal{T}$, which is compliant with an efficient usage of standard $\mathcal{LA}(\mathbb{Q})$-solvers/optimizers.

3.3. OMT($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$) wrt. other Optimization Problems

In this section we show that OMT($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$) captures many interesting optimizations problems.

LP is a particular subcase of OMT($\mathcal{LA}(\mathbb{Q})$) with no Boolean component, such that $\varphi \stackrel{\text{def}}{=} \varphi' \wedge (\mathsf{cost} = \sum_i \mathbf{a}_i \mathbf{x}_i)$ and $\varphi' = \bigwedge_j (\sum_i \mathbf{A}_{ij} \mathbf{x}_i \leq \mathbf{b}_j)$.

LDP can also be encoded into OMT($\mathcal{LA}(\mathbb{Q})$), since (2) and (3) can be written respectively as

$$\bigvee_i \bigwedge_j (\mathbf{A}^i_j \mathbf{x} \geq \mathbf{b}^i_j) \tag{8}$$

$$\bigwedge_j (\mathbf{A}_j \mathbf{x} \geq \mathbf{b}_j) \wedge \bigwedge_{j=1}^{t} \bigvee_{k \in I_j} (\mathbf{c}^k \mathbf{x} \geq d^k), \tag{9}$$

where $\mathbf{A}^i_j$ and $\mathbf{A}_j$ are respectively the jth row of the matrices $\mathbf{A}^i$ and $\mathbf{A}$, $\mathbf{b}^i_j$ and $\mathbf{b}_j$ are respectively the jth row of the vectors $\mathbf{b}^i$ and $\mathbf{b}$. Since (8) is not in CNF, the CNF-ization process of [Plaisted and Greenbaum 1986] is then applied.

LGDP (4) is straightforwardly encoded into a OMT($\mathcal{LA}(\mathbb{Q})$) problem $\langle \varphi, \mathsf{cost} \rangle$:

$$\begin{array}{rl} \varphi \stackrel{\text{def}}{=} & (\mathsf{cost} = \sum_{k \in K} \mathbf{z}_k + \mathbf{d}\mathbf{x}) \wedge [[\mathbf{B}\mathbf{x} \leq \mathbf{b}]] \wedge \phi \wedge [[\mathbf{0} \leq \mathbf{x}]] \wedge [[\mathbf{x} \leq \mathbf{e}]] \\ \wedge & \bigwedge_{k \in K} \bigvee_{j \in J_k} (Y^{jk} \wedge [[\mathbf{A}^{jk}\mathbf{x} \leq \mathbf{a}^{jk}]] \wedge (\mathbf{z}_k = c^{jk})) \\ \wedge & \bigwedge_{k \in K} ((\mathbf{z}_k \geq min_{j \in J_k} c^{jk}) \wedge (\mathbf{z}_k \leq max_{j \in J_k} c^{jk})) \end{array} \tag{10}$$

s.t. $[[\mathbf{x} \bowtie \mathbf{a}]]$ and $[[\mathbf{A}\mathbf{x} \bowtie \mathbf{a}]]$ are abbreviations respectively for $\bigwedge_i (\mathbf{x}_i \bowtie \mathbf{a}_i)$ and $\bigwedge_i (\mathbf{A}_{i\cdot}\mathbf{x} \bowtie \mathbf{a}_i)$, $\bowtie \in \{=, \neq \leq, \geq, <, >\}$. The last conjunction "$\bigwedge_{k \in K}((\mathbf{z}_k \geq \ \ldots\ ))$" in (10) is not necessary, but it improves the performances of the SMT($\mathcal{LA}(\mathbb{Q})$) solver, because it allows for exploiting the early-pruning SMT technique (see §2.2) by providing a range for the values of the $\mathbf{z}_k$'s before the respective Y^{jk}'s are assigned. Since (10) is not in CNF, the CNF-ization process of [Plaisted and Greenbaum 1986] is then applied.

Pseudo-Boolean (PB) constraints (see [Roussel and Manquinho 2009]) in the form $(\sum_i \mathbf{a}_i X^i \leq b)$ s.t. X^i are Boolean atoms and $\mathbf{a}_i$ constant values in $\mathbb{Q}$, and cost functions $\mathsf{cost} = \sum_i \mathbf{a}_i X^i$, are encoded into OMT($\mathcal{LA}(\mathbb{Q})$) by rewriting each PB-term $\sum_i \mathbf{a}_i X^i$ into the $\mathcal{LA}(\mathbb{Q})$-term $\sum_i \mathbf{x}_i$, $\mathbf{x}$ being an array of fresh $\mathcal{LA}(\mathbb{Q})$ variables, and by conjoining to φ the formula:

$$\textstyle\bigwedge_i ((\neg X^i \vee (\mathbf{x}_i = \mathbf{a}_i)) \wedge (X^i \vee (\mathbf{x}_i = 0)) \wedge (\mathbf{x}_i \geq 0) \wedge (\mathbf{x}_i \leq \mathbf{a}_i)\ ). \tag{11}$$

The term "$(\mathbf{x}_i \geq 0) \wedge (\mathbf{x}_i \leq \mathbf{a}_i)$" in (11) is not necessary, but it improves the performances of the SMT($\mathcal{LA}(\mathbb{Q})$) solver, because it allows for exploiting the early-pruning technique by providing a range for the values of the $\mathbf{x}_i$'s before the respective X^i's are assigned.

A (partial weighted) MaxSMT problem (see [Nieuwenhuis and Oliveras 2006; Cimatti et al. 2010; Cimatti et al. 2013a]) is a pair $\langle \varphi_h, \varphi_s \rangle$ where φ_h is a set of "hard" $\mathcal{T}$-clauses and φ_s is a set of weighted "soft" $\mathcal{T}$-clauses, s.t. a positive weight $\mathbf{a}_i$ is associated to each soft $\mathcal{T}$-clause $C_i \in \varphi_s$; the problem consists in finding a maximum-weight set of soft $\mathcal{T}$-clauses ψ_s s.t. $\psi_s \subseteq \varphi_s$ and $\varphi_h \cup \psi_s$ is $\mathcal{T}$-satisfiable. Notice that one can see $\mathbf{a}_i$ as a penalty to pay when the corresponding soft clause is not satisfied. A MaxSMT problem $\langle \varphi_h, \varphi_s \rangle$ can be encoded

straightforwardly into an SMT problem with PB cost function $\langle \varphi', \mathsf{cost} \rangle$ by augmenting each soft $\mathcal{T}$-clause C_j with a fresh Boolean variables X^j as follows:

$$\varphi' \stackrel{\text{def}}{=} \varphi_h \cup \textstyle\bigcup_{\mathcal{C}_j \in \varphi_s} \{(X^j \vee \mathcal{C}_j)\}; \;\; \mathsf{cost} \stackrel{\text{def}}{=} \sum_{\mathcal{C}_j \in \varphi_s} \mathbf{a}_j X^j. \tag{12}$$

Vice versa, $\langle \varphi', \mathsf{cost} \stackrel{\text{def}}{=} \sum_j \mathbf{a}_j X^j \rangle$ can be encoded into MaxSMT:

$$\varphi_h \stackrel{\text{def}}{=} \varphi'; \;\; \varphi_s \stackrel{\text{def}}{=} \textstyle\bigcup_j \{\underbrace{(\neg X^j)}_{\mathbf{a}_j}\}. \tag{13}$$

Thus, combining (11) and (12), optimization problems for SAT with PB constraints and MaxSAT can be encoded into OMT($\mathcal{LA}(\mathbb{Q})$), whilst those for SMT($\mathcal{T}$) with PB constraints and MaxSMT can be encoded into OMT($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$), under the assumption that $\mathcal{T}$ matches the definition in §3.1.

Remark 3.13. We notice the deep difference between OMT($\mathcal{LA}(\mathbb{Q})$)/OMT($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$) and the problem of SAT/SMT with PB constraints and cost functions (or MaxSAT/ MaxSMT) addressed in [Nieuwenhuis and Oliveras 2006; Cimatti et al. 2010; Cimatti et al. 2013a]. With the latter problems, the value of cost is a deterministic consequence of a truth assignment to the atoms of the formula, so that the search has only a Boolean component, consisting in finding the cheapest truth assignment. With OMT($\mathcal{LA}(\mathbb{Q})$)/ OMT($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$), instead, for every satisfying assignment μ it is also necessary to find the minimum-cost $\mathcal{LA}(\mathbb{Q})$-model for μ, so that the search has both a Boolean and a $\mathcal{LA}(\mathbb{Q})$-component.

4. PROCEDURES FOR OMT($\mathcal{LA}(\mathbb{Q})$) AND OMT($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$)

It might be noticed that very naive OMT($\mathcal{LA}(\mathbb{Q})$) or OMT($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$) procedures could be straightforwardly implemented by performing a sequence of calls to an SMT solver on formulas like $\varphi \wedge (\mathsf{cost} \geq \mathsf{l}_i) \wedge (\mathsf{cost} < \mathsf{u}_i)$, each time restricting the range $[\mathsf{l}_i, \mathsf{u}_i[$ according to a linear-search or binary-search schema. With the linear-search schema, every time the SMT solver returns a model of cost c_i, a new constraint $(\mathsf{cost} < c_i)$ would be added to φ, and the solver would be invoked again; however, the SMT solver would repeatedly generate the same $\mathcal{LA}(\mathbb{Q})$-satisfiable truth assignment, each time finding a cheaper model for it. With the binary-search schema the efficiency should improve; however, an initial lower-bound should be necessarily required as input (which is not the case, e.g., of the problems in §5.3.)

In this section we present more sophisticate procedures, based on the combination of SMT and minimization techniques. We first present and discuss an *offline* schema (§4.1) and an *inline* (§4.2) schema for an OMT($\mathcal{LA}(\mathbb{Q})$) procedure; then we show how to extend them to the OMT($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$) case (§4.3).

4.1. An offline schema for OMT($\mathcal{LA}(\mathbb{Q})$)

The general schema for the offline OMT($\mathcal{LA}(\mathbb{Q})$) procedure is displayed in Algorithm 1. It takes as input an instance of the OMT($\mathcal{LA}(\mathbb{Q})$) problem plus optionally values for lb and ub, which are implicitly considered to be $-\infty$ and $+\infty$ if not present, and returns the model $\mathcal{M}$ of minimum cost and its cost u (the value ub if φ is $\mathcal{LA}(\mathbb{Q})$-inconsistent). Notice that, by providing a lower bound lb [resp. an upper bound ub], the user implicitly assumes the responsibility of asserting there is no model whose cost is lower than lb [resp. there is a model whose cost is ub].

We represent φ as a set of clauses, which may be pushed or popped from the input formula-stack of an incremental SMT solver. To this extent, every operation like "$\varphi \leftarrow \varphi \cup \{...\}$" [resp. "$\varphi \leftarrow \varphi \setminus \{...\}$"] in Algorithm 1, "$\{...\}$" being a clause set, should be interpreted as "push $\{...\}$ into φ" [resp. "pop $\{...\}$ from φ"].

First, the variables l, u defining the current range are initialized to lb and ub respectively, the atom PIV to $\top$, and $\mathcal{M}$ is initialized to be an empty model. Then the procedure adds to φ the bound constraints, if present, which restrict the search within the range $[\mathsf{l}, \mathsf{u}[$ (row 2). (Obviously literals like $\neg(\mathsf{cost} < -\infty)$ and $(\mathsf{cost} < +\infty)$ are not added.) The solution space is then explored

Algorithm 1 Offline OMT($\mathcal{LA}(\mathbb{Q})$) Procedure based on Mixed Linear/Binary Search.

Require: $\langle \varphi, \mathsf{cost}, \mathsf{lb}, \mathsf{ub} \rangle$ {ub can be $+\infty$, lb can be $-\infty$}
1: $\mathsf{l} \leftarrow \mathsf{lb}$; $\mathsf{u} \leftarrow \mathsf{ub}$; $\mathsf{PIV} \leftarrow \top$; $\mathcal{M} \leftarrow \emptyset$
2: $\varphi \leftarrow \varphi \cup \{\neg(\mathsf{cost} < \mathsf{l}), (\mathsf{cost} < \mathsf{u})\}$ // Push bound constraints into φ
3: **while** ($\mathsf{l} < \mathsf{u}$) **do**
4: **if** (BinSearchMode()) **then** // Binary-search Mode
5: $\mathsf{pivot} \leftarrow \mathsf{ComputePivot}(\mathsf{l}, \mathsf{u})$
6: $\mathsf{PIV} \leftarrow (\mathsf{cost} < \mathsf{pivot})$
7: $\varphi \leftarrow \varphi \cup \{\mathsf{PIV}\}$ // Push PIV into φ
8: $\langle \mathsf{res}, \mu \rangle \leftarrow \mathsf{SMT.IncrementalSolve}(\varphi)$
9: **if** ($\mathsf{res} = \text{UNSAT}$) **then**
10: $\eta \leftarrow \mathsf{SMT.ExtractUnsatCore}(\varphi)$
11: **else**
12: $\eta \leftarrow \emptyset$
13: **end if**
14: **else** // Linear-search Mode
15: $\langle \mathsf{res}, \mu \rangle \leftarrow \mathsf{SMT.IncrementalSolve}(\varphi)$
16: $\eta \leftarrow \emptyset$
17: **end if**
18: **if** ($\mathsf{res} = \text{SAT}$) **then**
19: $\langle \mathcal{M}, \mathsf{u} \rangle \leftarrow \mathsf{Minimize}(\mathsf{cost}, \mu)$
20: $\varphi \leftarrow \varphi \cup \{(\mathsf{cost} < \mathsf{u})\}$ // Push new upper-bound constraint into φ
21: **else** // res = UNSAT
22: **if** ($\mathsf{PIV} \notin \eta$) **then**
23: $\mathsf{l} \leftarrow \mathsf{u}$
24: **else**
25: $\mathsf{l} \leftarrow \mathsf{pivot}$
26: $\varphi \leftarrow \varphi \setminus \{\mathsf{PIV}\}$ // Pop PIV from φ
27: $\varphi \leftarrow \varphi \cup \{\neg\mathsf{PIV}\}$ // Push $\neg$PIV into φ
28: **end if**
29: **end if**
30: **end while**
31: **return** $\langle \mathcal{M}, \mathsf{u} \rangle$

iteratively (rows 3-30), reducing the current range $[\mathsf{l}, \mathsf{u}[$ to explore at each loop, until the range is empty. Then $\langle \mathcal{M}, \mathsf{u} \rangle$ is returned —$\langle \emptyset, \mathsf{ub} \rangle$ if there is no solution in $[\mathsf{lb}, \mathsf{ub}[$— $\mathcal{M}$ being the model of minimum cost u. Each loop may work in either *linear-search* or *binary-search* mode, driven by the heuristic BinSearchMode(). Notice that if $\mathsf{u} = +\infty$ or $\mathsf{l} = -\infty$, then BinSearchMode() returns false.

In **linear-search mode**, steps 5-13 and 25-27 are not executed. First, an incremental SMT($\mathcal{LA}(\mathbb{Q})$) solver is invoked on φ (row 15). (Notice that, given the incrementality of the solver, every operation in the form "$\varphi \leftarrow \varphi \cup \{\phi_i\}$" [resp. $\varphi \leftarrow \varphi \setminus \{\phi_i\}$] is implemented as a "push" [resp. "pop"] operation on the stack representation of φ, see §2.1; it is also very important to recall that during the SMT call φ is updated with the clauses which are learned during the SMT search.) Then η is set to be empty, which forces condition 22 to hold.

If φ is $\mathcal{LA}(\mathbb{Q})$-satisfiable, then it is returned res =SAT and a $\mathcal{LA}(\mathbb{Q})$-satisfiable truth assignment μ for φ. Thus Minimize is invoked on (the subset of $\mathcal{LA}(\mathbb{Q})$-literals of) μ, [3] returning the model $\mathcal{M}$ for μ of minimum cost u ($-\infty$ iff the problem is unbounded). The current solution u becomes the new upper bound, thus the $\mathcal{LA}(\mathbb{Q})$-atom $(\mathsf{cost} < \mathsf{u})$ is added to φ (row 20). Notice that, if the

[3] Possibly after applying pure-literal filtering to μ (see §2.2).

problem is unbounded, then for some μ Minimize will return $-\infty$, forcing condition 3 to be false and the whole process to stop. If φ is $\mathcal{LA}(\mathbb{Q})$-unsatisfiable, then no model in the current cost range $[\mathsf{l}, \mathsf{u}[$ can be found; hence the flag l is set to u, forcing the end of the loop.

In **binary-search mode** at the beginning of the loop (steps 5-13), the value $\mathsf{pivot} \in]\mathsf{l}, \mathsf{u}[$ is computed by the heuristic function ComputePivot (in the simplest form, pivot is $(\mathsf{l} + \mathsf{u})/2$), the possibly-new atom $\mathsf{PIV} \stackrel{\text{def}}{=} (\mathsf{cost} < \mathsf{pivot})$ is pushed into the formula stack, so that to temporarily restrict the cost range to $[\mathsf{l}, \mathsf{pivot}[$. Then the incremental SMT solver is invoked on φ; if the result is UNSAT, the feature SMT.ExtractUnsatCore is activated, which returns also a subset η of formulas in (the formula stack of) φ which caused the unsatisfiability of φ (see §2.1). This exploits techniques similar to unsat-core extraction [Lynce and Marques-Silva 2004]. (In practical implementations, it is not strictly necessary to explicitly produce the unsat core η; rather, it suffices to check if $\mathsf{PIV} \in \eta$.)

If φ is $\mathcal{LA}(\mathbb{Q})$-satisfiable, then the procedure behaves as in linear-search mode. If instead φ is $\mathcal{LA}(\mathbb{Q})$-unsatisfiable, we look at η and distinguish two subcases. If PIV does not occur in η, this means that $\varphi \setminus \{\mathsf{PIV}\}$ is $\mathcal{LA}(\mathbb{Q})$-inconsistent, i.e. there is no model in the whole cost range $[\mathsf{l}, \mathsf{u}[$. Then the procedure behaves as in linear-search mode, forcing the end of the loop. Otherwise, we can only conclude that there is no model in the cost range $[\mathsf{l}, \mathsf{pivot}[$, so that we still need exploring the cost range $[\mathsf{pivot}, \mathsf{u}[$. Thus l is set to pivot, PIV is popped from φ and its negation is pushed into φ. Then the search proceeds, investigating the cost range $[\mathsf{pivot}, \mathsf{u}[$.

We notice an important fact: if BinSearchMode() always returned true, then Algorithm 1 would not necessarily terminate. In fact, an SMT solver invoked on φ may return a set η containing PIV even if $\varphi \setminus \mathsf{PIV}$ is $\mathcal{LA}(\mathbb{Q})$-inconsistent. [4] Thus, e.g., the procedure might got stuck into a "Zeno" [5] infinite loop, each time halving the cost range right-bound (e.g., $[-1, 0[$, $[-1/2, 0[$, $[-1/4, 0[$,...).

To cope with this fact, however, it suffices to guarantee that BinSearchMode() returns false after a finite number of such steps, guaranteeing thus that eventually a linear-search loop will be forced, which detects the inconsistency. In our implementations, we have empirically chosen to force one linear-search loop after *every* binary-search loop returning UNSAT, because satisfiable calls are typically much cheaper than unsatisfiable ones. We have empirically verified in previous tests that this was in general the best option, since introducing this test caused no significant overhead and prevented the chains of (very expensive) unsatisfiable calls where they used to occur.

Under such hypothesis, as a consequence of Theorem 3.6 of §3.2, we have that:

(i) Algorithm 1 terminates. In linear-search mode it terminates because there are only a finite number of candidate truth assignments μ to be enumerated, and steps 19-20 guarantee that the same assignment μ will never be returned twice by the SMT solver. In mixed linear/binary-search mode, as above, it terminates since there can be at most finitely-many binary-search loops between two consequent linear-search loops;

(ii) Algorithm 1 returns a model of minimum cost, because it explores the whole search space of candidate truth assignments, and for every suitable assignment μ Minimize finds the minimum-cost model for μ;

(iii) Algorithm 1 requires polynomial space, under the assumption that the underlying CDCL SAT solver adopts a polynomial-size clause-deleting strategy (which is typically the case of SMT solvers, including MATHSAT).

4.1.1. Handling strict inequalities. Minimize is a simple extension of the simplex-based $\mathcal{LA}(\mathbb{Q})$-Solver of [Dutertre and de Moura 2006], which is invoked after one solution is found, minimizing it by standard simplex techniques. We recall that the algorithm in [Dutertre and de Moura 2006] can

[4] A CDCL-based SMT solver implicitly builds a resolution refutation whose leaves are either clauses in φ or $\mathcal{LA}(\mathbb{Q})$-lemmas, and the set η represents the subset of clauses in φ which occur as leaves of such proof (see e.g. [Cimatti et al. 2011] for details). If the SMT solver is invoked on φ even $\varphi \setminus \mathsf{PIV}$ is $\mathcal{LA}(\mathbb{Q})$-inconsistent, then it can "use" PIV and return a proof involving it even though another PIV-less proof exists.

[5] In the famous Zeno's paradox, Achilles never reaches the tortoise for a similar reason.

handle strict inequalities. Thus, if μ contains strict inequalities, then Minimize temporarily relaxes them into non-strict ones and then it finds a solution $\mathbf{x}$ of minimum cost min of the relaxed problem, namely μ_{rel}. (Notice that this could be done also by any standard LP package.) Then:

(1) if such minimum-cost solution $\mathbf{x}$ lays only on non-strict inequalities, then $\mathbf{x}$ is a also solution of the non-relaxed problem μ, hence min can be returned;
(2) otherwise, we temporarily add the constraint $(\mathsf{cost} \leq \mathsf{min})$ to the non-relaxed version of μ and then we invoke on it the $\mathcal{LA}(\mathbb{Q})$-solving procedure of [Dutertre and de Moura 2006] (without minimization), since such algorithm can handle strict inequalities. Then:
   (i) if the procedure returns SAT, then μ has a model of cost min. If so, then the value min can be returned, and $(\mathsf{cost} < \mathsf{min})$ can be pushed into φ;
   (ii) otherwise, μ has no model of cost min. If so, since μ has a convex set of solutions whose cost is strictly greater than min and there is a solution of cost min for the relaxed problem μ_{rel}, then for some $\delta > 0$ and for every cost $c \in]\mathsf{min}, \mathsf{min} + \delta]$ there exists a solution for μ of cost c. (If needed explicitly, such solution can be computed using the techniques for handling strict inequalities described in [Dutertre and de Moura 2006].) Thus the value min can be tagged as a non-strict minimum and returned, so that the constraint $(\mathsf{cost} \leq \mathsf{min})$, rather than $(\mathsf{cost} < \mathsf{min})$, is pushed into φ.

Notice that situation (2).(i) is very rare in practice but it is possible in principle, as illustrated in the following example.

Example 4.1. Suppose we have that $\mu = \{(\mathsf{cost} \geq 1), (\mathsf{cost} > y), (\mathsf{cost} > -y)\}$. If we temporarily relax strict inequalities into non-strict ones, then $\{\mathsf{cost} = 1, y = 1\}$ is a minimum-cost solution which lays on the strict inequality $(\mathsf{cost} > y)$. Nevertheless, there is a solution of cost 1 for the un-relaxed problem, e.g., $\{\mathsf{cost} = 1, y = 0.9999\}$.

Notice also that $(\mathsf{cost} \leq \mathsf{min})$ is pushed into φ only if the minimum cost of the current assignment μ is strictly greater than min, as in situation (2).(ii). This prevents the SMT solver from returning μ again. Therefore the termination, the correctness and the completeness of the algorithm are guaranteed also in the case some truth assignments have strict minimum costs.

4.1.2. Discussion. We remark a few facts about this procedure.

(1) If Algorithm 1 is interrupted (e.g., by a timeout device), then u can be returned, representing the best approximation of the minimum cost found so far.
(2) The incrementality of the SMT solver (see §2.1 and §2.2) plays an essential role here, since at every call $\mathsf{SMT.IncrementalSolve}$ resumes the status of the search at the end of the previous call, only with tighter cost range constraints. (Notice that at each call here the solver can reuse all previously-learned clauses.) To this extent, one can see the whole process mostly as only one SMT process, which is interrupted and resumed each time a new model is found, in which cost range constraints are progressively tightened.
(3) In Algorithm 1 all the literals constraining the cost range (i.e., $\neg(\mathsf{cost} < \mathsf{l})$, $(\mathsf{cost} < \mathsf{u})$) are added to φ as unit clauses; thus inside $\mathsf{SMT.IncrementalSolve}$ they are immediately unit-propagated, becoming part of each truth assignment μ from the very beginning of its construction. As soon as novel $\mathcal{LA}(\mathbb{Q})$-literals are added to μ which prevent it from having a $\mathcal{LA}(\mathbb{Q})$-model of cost in $[\mathsf{l}, \mathsf{u}[$, the $\mathcal{LA}(\mathbb{Q})$-solver invoked on μ by early-pruning calls (see §2.2) returns UNSAT and the $\mathcal{LA}(\mathbb{Q})$-lemma $\neg\eta$ describing the conflict $\eta \subseteq \mu$, triggering theory-backjumping and -learning. To this extent, $\mathsf{SMT.IncrementalSolve}$ implicitly plays a form of *branch & bound*: (i) decide a new literal l and unit- or theory-propagate the literals which derive from l ("branch") and (ii) backtrack as soon as the current branch can no more be expanded into models in the current cost range ("bound").
(4) The unit clause $\neg(\mathsf{cost} < \mathsf{l})$ plays a role even in linear-search mode, since it helps pruning the search inside $\mathsf{SMT.IncrementalSolve}$.

(5) In binary-search mode, the range-partition strategy may be even more aggressive than that of standard binary search, because the minimum cost u returned in row 19 can be smaller than pivot, so that the cost range is more then halved.
(6) Unlike with other domains (e.g., search in sorted arrays) here binary search is not "obviously faster" than linear search, because the unsatisfiable calls to SMT.IncrementalSolve are typically much more expensive than the satisfiable ones, since they must explore the whole Boolean search space rather than only a portion of it —although with a higher pruning power, due to the stronger constraint induced by the presence of pivot. Thus, we have a tradeoff between a typically much-smaller number of calls plus a stronger pruning power in binary search versus an average much smaller cost of the calls in linear search. To this extent, it is possible in principle to use dynamic/adaptive strategies for ComputePivot (see [Sellmann and Kadioglu 2008]).

4.2. An inline schema for OMT($\mathcal{LA}(\mathbb{Q})$)

With the inline schema, the whole optimization procedure is pushed inside the SMT solver by embedding the range-minimization loop inside the CDCL Boolean-search loop of the standard lazy SMT schema of §2.2. The SMT solver, which is thus called only once, is modified as follows.

Initialization. The variables $\mathsf{lb}, \mathsf{ub}, \mathsf{l}, \mathsf{u}, \mathsf{PIV}, \mathsf{pivot}, \mathcal{M}$ are brought inside the SMT solver, and are initialized as in Algorithm 1, steps 1-2.

Range Updating & Pivoting. Every time the search of the CDCL SAT solver gets back to decision level 0, the range $[\mathsf{l}, \mathsf{u}[$ is updated s.t. u [resp. l] is assigned the lowest [resp. highest] value u_i [resp. l_i] such that the atom $(\mathsf{cost} < \mathsf{u}_i)$ [resp. $\neg(\mathsf{cost} < \mathsf{l}_i)$] is currently assigned at level 0. (If $\mathsf{u} \leq \mathsf{l}$, or two literals $l, \neg l$ are both assigned at level 0, then the procedure terminates, returning the current value of u.) Then BinSearchMode() is invoked: if it returns true, then ComputePivot computes $\mathsf{pivot} \in]\mathsf{l}, \mathsf{u}[$, and the (possibly new) atom $\mathsf{PIV} \stackrel{\text{def}}{=} (\mathsf{cost} < \mathsf{pivot})$ is decided to be true (level 1) by the SAT solver. This mimics steps 5-7 in Algorithm 1, temporarily restricting the cost range to $[\mathsf{l}, \mathsf{pivot}[$.

Decreasing the Upper Bound. When an assignment μ propositionally satisfying φ is generated which is found $\mathcal{LA}(\mathbb{Q})$-consistent by $\mathcal{LA}(\mathbb{Q})$-Solver, μ is also fed to Minimize, returning the minimum cost min of μ; then the unit clause $(\mathsf{cost} < \mathsf{min})$ is learned and fed to the backjumping mechanism, which forces the SAT solver to backjump to level 0 and then to unit-propagate $(\mathsf{cost} < \mathsf{min})$. This case mirrors steps 18-20 in Algorithm 1, permanently restricting the cost range to $[\mathsf{l}, \mathsf{min}[$. Minimize is embedded within the $\mathcal{LA}(\mathbb{Q})$-Solver, so that it is called incrementally after it, without restarting its search from scratch.

As a result of these modifications, we also have the following typical scenario (see Figure 1).

Increasing the Lower Bound. In binary-search mode, when a conflict occurs s.t. the conflict analysis of the SAT solver produces a conflict clause in the form $\neg\mathsf{PIV} \vee \neg\eta'$ s.t. all literals in η' are assigned true at level 0 (i.e., $\varphi \wedge \mathsf{PIV}$ is $\mathcal{LA}(\mathbb{Q})$-inconsistent), then the SAT solver backtracks to level 0, unit-propagating $\neg\mathsf{PIV}$. This case mirrors steps 25-27 in Algorithm 1, permanently restricting the cost range to $[\mathsf{pivot}, \mathsf{u}[$.

Although the modified SMT solver mimics to some extent the behaviour of Algorithm 1, the "control" of the range-restriction process is handled by the standard SMT search. To this extent, notice that also other situations may allow for restricting the cost range: e.g., if $\varphi \wedge \neg(\mathsf{cost} < \mathsf{l}) \wedge (\mathsf{cost} < \mathsf{u}) \models (\mathsf{cost} \bowtie \mathsf{m})$ for some atom $(\mathsf{cost} \bowtie \mathsf{m})$ occurring in φ s.t. $\mathsf{m} \in [\mathsf{l}, \mathsf{u}[$ and $\bowtie \in \{\leq, <, \geq, >\}$, then the SMT solver may backjump to decision level 0 and propagate $(\mathsf{cost} \bowtie \mathsf{m})$, further restricting the cost range.

The same facts (1)-(6) about the offline procedure in §4.1 hold for the inline version. The efficiency of the inline procedure can be further improved as follows.

Activating previously-learned clauses. In binary-search mode, when an assignment μ with a novel minimum min is found, not only $(\mathsf{cost} < \mathsf{min})$ but also $\mathsf{PIV} \stackrel{\text{def}}{=} (\mathsf{cost} < \mathsf{pivot})$ is learned as unit clause, although the latter is redundant from the logical perspective because $\mathsf{min} < \mathsf{pivot}$. In fact, the unit clause PIV allows the SAT solver for reusing all the clauses in the form $\neg\mathsf{PIV} \vee C_i$ which have

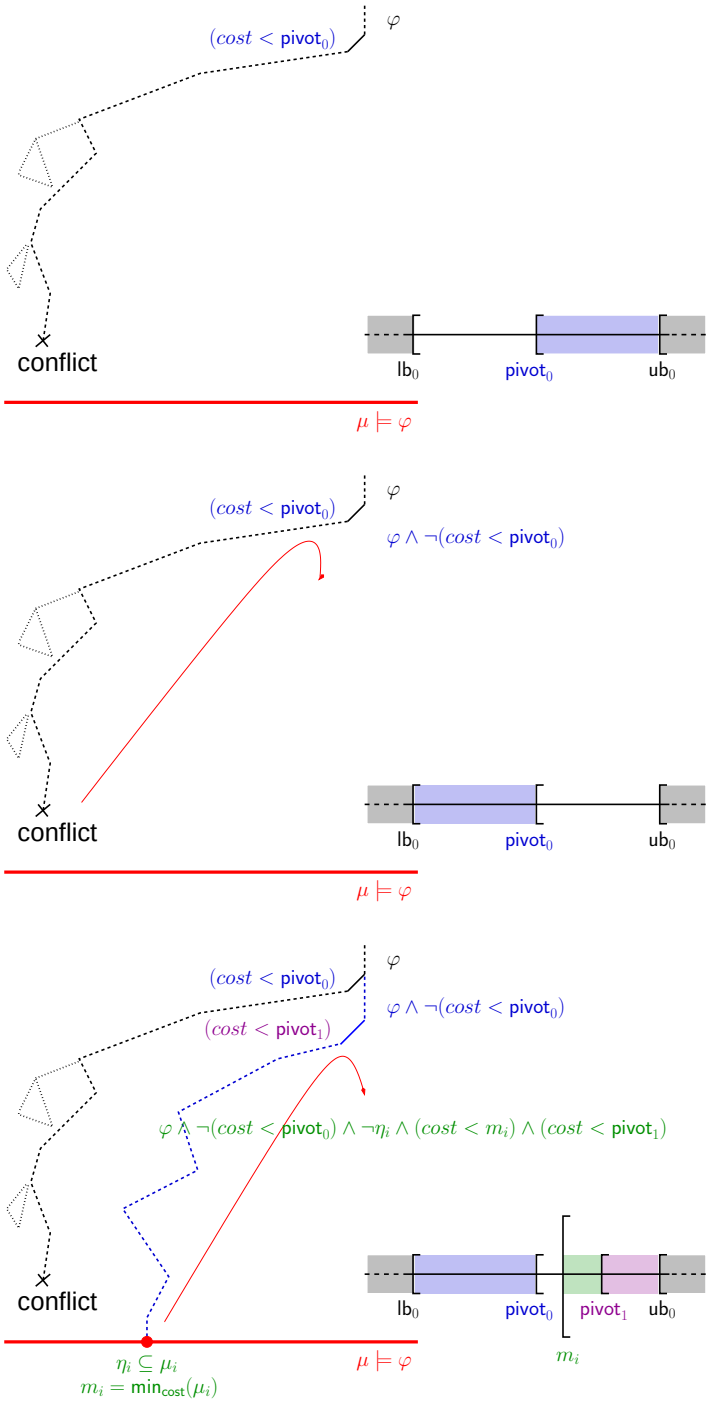

Fig. 1. One piece of possible execution of an inline procedure. (i) Pivoting on $(\mathsf{cost} < \mathsf{pivot}_0)$. (ii) Increasing the lower bound to pivot_0. (iii) Decreasing the upper bound to $\mathsf{min}_{\mathsf{cost}}(\mu_i)$.

been learned when investigating the cost range $[\mathsf{l}, \mathsf{pivot}[$, by unit-resolving them into the corresponding clauses C_i.. (In Algorithm 1 this is done implicitly, since PIV is not popped from φ before step 26.) Notice that the above trick is useful because the algorithm of [Dutertre and de Moura 2006] is not "$\mathcal{T}$-deduction-complete", that is, it is not guaranteed to $\mathcal{T}$-deduce PIV from $\{..., (\mathsf{cost} < \mathsf{min})\}$.

In addition, the $\mathcal{LA}(\mathbb{Q})$-inconsistent assignment $\mu \wedge (\mathsf{cost} < \mathsf{min})$ may be fed to $\mathcal{LA}(\mathbb{Q})$-Solver and the negation of the returned conflict $\neg\eta \vee \neg(\mathsf{cost} < \mathsf{min})$ s.t. $\eta \subseteq \mu$, can be learned, preventing the SAT solver from generating any assignment containing η in the future.

Tightening. In binary-search mode, if $\mathcal{LA}(\mathbb{Q})$-Solver returns a conflict set $\eta \cup \{\mathsf{PIV}\}$, then it is further asked to find the maximum value max s.t. $\eta \cup \{(\mathsf{cost} < \mathsf{max})\}$ is $\mathcal{LA}(\mathbb{Q})$-inconsistent. (This is done with a simple modification of the algorithm in [Dutertre and de Moura 2006].)

— If $\mathsf{max} \geq \mathsf{u}$, then the clause $C^* \stackrel{\text{def}}{=} \neg\eta \vee \neg(\mathsf{cost} < \mathsf{u})$ is used do drive backjumping and learning instead of $C \stackrel{\text{def}}{=} \neg\eta \vee \neg\mathsf{PIV}$. Since the unit clause $(\mathsf{cost} < \mathsf{u})$ is permanently assigned at level 0, this is equivalent to learning only $\neg\eta$, so that the dependency of the conflict from PIV is removed. Eventually, instead of using C to drive backjumping to level 0 and then to propagate $\neg\mathsf{PIV}$, the SMT solver may use C^* (which is the same as using $\neg\eta$), then forcing the procedure to stop.

— If $\mathsf{u} > \mathsf{max} > \mathsf{pivot}$, then the clauses $C_1 \stackrel{\text{def}}{=} \neg\eta \vee \neg(\mathsf{cost} < \mathsf{max})$ and $C_2 \stackrel{\text{def}}{=} \neg\mathsf{PIV} \vee (\mathsf{cost} < \mathsf{max})$ are used to drive backjumping and learning instead of $C \stackrel{\text{def}}{=} \neg\eta \vee \neg\mathsf{PIV}$. (Notice that C can be inferred by resolving C_1 and C_2.) In particular, C_2 forces backjumping to level 1 and unit-propagating the (possibly fresh) atom $(\mathsf{cost} < \mathsf{max})$; eventually, instead of using C do drive backjumping to level 0 and then to propagate $\neg\mathsf{PIV}$, the SMT solver may use C_1 for backjumping to level 0 and then to propagate $\neg(\mathsf{cost} < \mathsf{max})$, restricting the range to $[\mathsf{max}, \mathsf{u}[$ rather than to $[\mathsf{pivot}, \mathsf{u}[$.

Notice that tightening is useful because the algorithm of [Dutertre and de Moura 2006] is guaranteed neither to find the "tightest" theory conflict $\eta \cup \{(\mathsf{cost} < \mathsf{max})\}$, nor to $\mathcal{T}$-deduce $(\mathsf{cost} < \mathsf{max})$ from $\{..., \mathsf{PIV}\}$.

Example 4.2. Consider the formula $\varphi \stackrel{\text{def}}{=} \psi \wedge (\mathsf{cost} \geq a + 15) \wedge (a \geq 0)$ for some ψ in the cost range $[0, 16[$. With binary-search deciding $\mathsf{PIV} \stackrel{\text{def}}{=} (\mathsf{cost} < 8)$, the $\mathcal{LA}(\mathbb{Q})$-Solver produces the lemma $C \stackrel{\text{def}}{=} \neg(\mathsf{cost} \geq a + 15) \vee \neg(a \geq 0) \vee \neg\mathsf{PIV}$, causing a backjumping step to level 0 on C and the unit-propagation of $\neg\mathsf{PIV}$, restricting the range to $[8, 16[$; it takes a sequence of similar steps to progressively restrict the range to $[12, 16[$, $[14, 16[$, and $[15, 16[$. If instead the $\mathcal{LA}(\mathbb{Q})$-Solver produces the lemmas $C_1 \stackrel{\text{def}}{=} \neg(\mathsf{cost} \geq a + 15) \vee \neg(a \geq 0) \vee \neg(\mathsf{cost} < 15)$ and $C_2 \stackrel{\text{def}}{=} \neg\mathsf{PIV} \vee (\mathsf{cost} < 15)$, then this first causes a backjumping step on C_2 to level 1 with the unit-propagation of $(\mathsf{cost} < 15)$, and then a backjumping step on C_1 to level zero with the unit-propagation of $\neg(\mathsf{cost} < 15)$, which directly restricts the range to $[15, 16[$.

Adaptive Mixed Linear/Binary Search Strategy. An adaptive version of the heuristic $\mathsf{BinSearchMode}()$ decides the next search mode according to the ratio between the progress obtained in the latest binary- and linear-search steps and their respective costs. If either ub or lb is not present —or if we are immediately after an UNSAT binary-search step, in compliance with the strategy to avoid infinite "Zeno" sequences described in §4.1— then the heuristic selects linear-search mode. Otherwise, it selects binary-search mode if and only if

$$\left|\frac{\Delta\mathsf{ub}_{lin}}{\Delta\#\mathsf{conf}_{lin}}\right| < \left|\frac{\Delta\mathsf{ub}_{bin}}{\Delta\#\mathsf{conf}_{bin}}\right|,$$

where $\Delta\mathsf{ub}_{lin}$ and $\Delta\mathsf{ub}_{bin}$ are respectively the variations of the upper bound ub in the latest linear-search and SAT binary-search steps performed, estimating the progress achieved by such steps, whilst $\Delta\#\mathsf{conf}_{lin}$ and $\Delta\#\mathsf{conf}_{bin}$ are respectively the number of conflicts produced in such steps, estimating their expense.

Overall, the inline version described in this section presents some potential computational advantages wrt. the offline version of Algorithm 1. First, despite the incrementality of the calls to the SMT solver, suspending and resuming it may cause some overhead, because at every call the decision stack is popped to decision level 0, so that some extra decisions, unit-propagations and early-pruning calls to the $\mathcal{T}$-Solver may be necessary to get back to the previous search status. Second, in Algorithm 1 the procedure Minimize is invoked from scratch in a non-incremental way, whilst in the inline version it is embedded inside the $\mathcal{LA}(\mathbb{Q})$-Solver, so that it starts the minimization process from an existing solution rather than from scratch. Third, Algorithm 1 requires computing the unsatisfiable core of φ—or at least checking if PIV belongs to such unsat core— which causes overhead. Notice that the problem of computing efficiently minimal unsat-cores in SMT is still ongoing research (see [Cimatti et al. 2011]), so that in Algorithm 1 there is a tradeoff between the cost of reducing the size of the cores and the probability of performing useless optimization steps.

4.3. Extensions to OMT$(\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T})$

We recall the terminology, assumptions, definitions and results of §3.2. Theorems 3.6, 3.10 and 3.12 allow for extending to the OMT$(\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T})$ case the procedures of §4.1 and §4.2 as follows.

As suggested by Theorem 3.10, straightforward OMT$(\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T})$ extensions of the procedures for OMT$(\mathcal{LA}(\mathbb{Q}))$ of §4.1 and §4.2 would be such that the SMT solver enumerates ed-extended satisfying truth assignments $\eta \stackrel{\text{def}}{=} \mu \wedge \mu_{ed}$ as in Definition 3.9, checking the $\mathcal{T}$- and $\mathcal{LA}(\mathbb{Q})$-consistency of its components $\mu_{\mathcal{T}} \wedge \mu_{ed}$ and $\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ed}$ respectively, and then minimizing the $\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ed}$ component. Termination is guaranteed by the fact that each $\mathcal{EX}_{ed}(\mu)$ is a finite set, whilst correctness and completeness is guaranteed by Theorems 3.6 and 3.10.

Nevertheless, as suggested in §3.2, minimizing $\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ed}$ efficiently could be problematic due to the presence of negated interface equalities $\neg(x_i = x_j)$. Thus, alternative "asymmetric" procedures, in compliance with the efficient usage of $\mathcal{LA}(\mathbb{Q})$-solvers in SMT, should instead enumerate edi-extended satisfying truth assignments $\rho \stackrel{\text{def}}{=} \mu \wedge \mu_{eid}$ as in Definition 3.11, checking the $\mathcal{T}$- and $\mathcal{LA}(\mathbb{Q})$-consistency of its components $\mu_{\mathcal{T}} \wedge \mu_{ed}$ and $\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ei}$ respectively, and then minimizing the $\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ei}$ component. This prevents from passing negated interface equalities to Minimize. As before, termination is guaranteed by the fact that each $\mathcal{EX}_{edi}(\mu)$ is a finite set, whilst correctness and completeness is guaranteed by Theorems 3.6 and 3.12.

This motivates and explains the following OMT$(\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T})$ variants of the offline and inline procedures of §4.1 and §4.2 respectively.

Algorithm 1 is modified as follows. First, SMT.IncrementalSolve in steps 8 and 15 is asked to return also a $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-model $\mathcal{I}$. Then in step 19 Minimize is invoked instead on $\langle \mathsf{cost}, \mu_{\mathcal{LA}(\mathbb{Q})} \cup \mu_{ei} \rangle$, s.t.

$$\begin{array}{rll} \mu_{ei} \stackrel{\text{def}}{=} & \{(x_i = x_j), \neg(x_i < x_j), \neg(x_i > x_j) & \mid (x_i = x_j) \in \mathcal{IE}(\mu),\ \mathcal{I} \models (x_i = x_j)\} \\ \cup & \{(x_i < x_j), \neg(x_i > x_j) & \mid (x_i = x_j) \in \mathcal{IE}(\mu),\ \mathcal{I} \models (x_i < x_j)\} \\ \cup & \{(x_i > x_j), \neg(x_i < x_j) & \mid (x_i = x_j) \in \mathcal{IE}(\mu),\ \mathcal{I} \models (x_i > x_j)\}. \end{array}$$

In practice, the negated strict inequalities $\neg(x_i < x_j), \neg(x_i > x_j)$ are omitted from μ_{ei}, because they are entailed by the corresponding non-negated equalities/inequalities.

The implementation of an inline OMT$(\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T})$ procedures comes nearly for free once the SMT solver handles $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-solving by *Delayed Theory Combination* [Bozzano et al. 2006], with the strategy of case-splitting automatically disequalities $\neg(x_i = x_j)$ into the two inequalities $(x_i < x_j)$ and $(x_i > x_j)$, which is implemented in MATHSAT: the solver enumerates truth assignments in the form $\rho \stackrel{\text{def}}{=} \mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{eid} \wedge \mu_{\mathcal{T}}$ as in Definition 3.11, and passes $\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ei}$ and $\mu_{\mathcal{T}} \wedge \mu_{ed}$ to the $\mathcal{LA}(\mathbb{Q})$-Solver and $\mathcal{T}$-Solver respectively. (Notice that this strategy, although not explicitly described in [Bozzano et al. 2006], implicitly implements points (a) and (c) of Theorem 3.12.) If so, then, in accordance with points (b) and (d) of Theorem 3.12, it suffices to apply Minimize to $\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ei}$, then learn $(\mathsf{cost} < \mathsf{min})$ and use it for backjumping, as in §4.2. As with

the offline version, in practice the negated strict inequalities are omitted from μ_{ei}, because they are entailed by the corresponding non-negated equalities/inequalities.

5. EXPERIMENTAL EVALUATION

We have implemented both the offline and inline OMT($\mathcal{LA}(\mathbb{Q})$) procedures and the inline OMT($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$) procedures of §4 on top of MATHSAT5 [6] [Cimatti et al. 2013b]; we refer to them as OPTIMATHSAT. MATHSAT5 is a state-of-the-art SMT solver which supports most of the quantifier-free SMT-LIB theories and their combinations, and provides many other SMT functionalities (like, e.g., unsat-core extraction [Cimatti et al. 2011], interpolation [Cimatti et al. 2010], All-SMT [Cavada et al. 2007]).

We consider different configurations of OPTIMATHSAT, depending on the approach (offline vs. inline, denoted by "-OF" and "-IN") and on the search schema (linear vs. binary vs. adaptive, denoted respectively by "-LIN", "-BIN" and "-ADA"). [7] For example, the configuration OPTIMATHSAT-LIN-IN denotes the inline linear-search procedure. We used only five configurations since the "-ADA-OF" were not implemented.

Due to the absence of competitors on OMT($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$), we evaluate the performance of our five configurations of OPTIMATHSAT by comparing them against the commercial LGDP tool GAMS [8] v23.7.1 [Brooke et al. 2011] on OMT($\mathcal{LA}(\mathbb{Q})$) problems. GAMS is a tool for modeling and solving optimization problems, consisting of different language compilers, which translate mathematical problems into representations required by specific solvers, like CPLEX [IBM 2010]. GAMS provides two reformulation tools, LOGMIP [9] v2.0 and JAMS [10] (a new version of the EMP [11] solver), s.t. both of them allow for reformulating LGDP models by using either big-M (BM) or convex-hull (CH) methods [Raman and Grossmann 1994; Sawaya and Grossmann 2012]. We use CPLEX v12.2 [IBM 2010] (through an OSI/CPLEX link) to solve the reformulated MILP models. All the tools were executed using default options, as suggested by the authors [Vecchietti 2011]. We also compared OPTIMATHSAT against MATHSAT augmented by Pseudo-Boolean (PB) optimization [Cimatti et al. 2010] (we call it PB-MATHSAT) on MaxSMT problems.

Remark 5.1. Importantly, MATHSAT and OPTIMATHSAT use *infinite-precision arithmetic* whilst the GAMS tools and CPLEX implement standard *floating-point arithmetic*. Moreover the former handle strict inequalities natively (see §2.2), whilst the GAMS tools use an approximation with a very-small constant value "`eps`" ϵ (default $\epsilon \stackrel{\text{def}}{=} 10^{-6}$), so that, e.g., "$(x > 0)$ *is internally rewritten into* $(x \geq 10^{-6})$" [12].

The comparison is run on four distinct collections of benchmark problems:

— (§5.2) LGDP problems, proposed by LOGMIP and JAMS authors [Vecchietti and Grossmann 2004; Sawaya and Grossmann 2005; Sawaya and Grossmann 2012];
— (§5.3) OMT($\mathcal{LA}(\mathbb{Q})$) problems from SMT-LIB [13];
— (§5.4) OMT($\mathcal{LA}(\mathbb{Q})$) problems, coming from encoding parametric verification problems from the SAL [14] model checker;
— (§5.5) the MaxSMT problems from [Cimatti et al. 2010].

The encodings from LGDP to OMT($\mathcal{LA}(\mathbb{Q})$) and back are described in §5.1.

[6] http://mathsat.fbk.eu/.
[7] Here "-LIN" means that BinSearchMode() always returns false, "-BIN" denotes the mixed linear-binary strategy described in §4.1 to ensure termination, whilst "-ADA" refers to the adaptive strategy illustrated in §4.2.
[8] http://www.gams.com.
[9] http://www.logmip.ceride.gov.ar/index.html.
[10] http://www.gams.com/.
[11] http://www.gams.com/dd/docs/solvers/emp.pdf.
[12] GAMS support team, email personal communication, 2012.
[13] http://www.smtlib.org/.
[14] http://sal.csl.sri.com.

All tests were executed on two identical 2.66 GHz Xeon machines with 4GB RAM running Linux, using a timeout of 600 seconds for each run. In order to have a reliable and fair measurement of CPU time, we have run only one process per PC at a time. Overall, the evaluation consisted in $\approx 40,000$ solver runs, for a total CPU time of up to 276 CPU days.

The correctness of the minimum costs min found by OPTIMATHSAT have been cross-checked by another SMT solver, YICES [15] by checking the inconsistency within the bounds of $\varphi \wedge (\mathsf{cost} < \mathsf{min})$ and the consistency of $\varphi \wedge (\mathsf{cost} = \mathsf{min})$ (if min is non-strict), or of $\varphi \wedge (\mathsf{cost} \leq \mathsf{min})$ and $\varphi \wedge (\mathsf{cost} = \mathsf{min} + \epsilon)$ (if min is strict), ϵ being some very small value.

All versions of OPTIMATHSAT passed the above checks. On the LGDP problems (§5.2) all tools agreed on the final results, apart from tiny rounding errors by GAMS tools; [16] on all the other problem collections (§5.3, §5.4, §5.5) instead, the results of the GAMS tools were affected by errors, which we will discuss there.

In order to make the experiments reproducible, more detailed tables, the full-size plots, a Linux binary of OPTIMATHSAT, the problems, and the results are made available. [17] (We cannot distribute the GAMS tools since they are subject to licencing restrictions, see [Brooke et al. 2011]; however, they can be obtained at GAMS url.)

5.1. Encodings.

In order to translate LGDP models into OMT($\mathcal{LA}(\mathbb{Q})$) problems we use the encoding in (10) of §3.3, namely LGDP2SMT. Notice that LGDP models are written in GAMS language which provides a large number of constructs. Since our encoder supports only base constructs (like equations and disjunctions), before generating the LGDP2SMT encoding, we used the GAMS Converter tool for converting complex GAMS specifications (e.g. containing sets and indexed equations) into simpler specifications. Notice also that in the GAMS language the disjunction of constraints in (10) must be described as nested if-then-elses on the Boolean propositions Y^{jk}, so that to avoid the need of including explicitly in ϕ the "xor" constraints discussed in the explanation of (10). Our encodings in both directions comply with this fact.

In order to translate OMT($\mathcal{LA}(\mathbb{Q})$) problems into LGDP models we consider two different encodings, namely SMT2LGDP$_1$ and SMT2LGDP$_2$.

Since GAMS tools do not handle negated equalities and strict inequalities, with both encodings negated equalities $\neg(t_1 = t_2)$ or $(t_1 \neq t_2)$ in the input $\mathcal{LA}(\mathbb{Q})$-formula φ are first replaced by the disjunction of two inequalities $\neg(t_1 \leq t_2) \vee \neg(t_1 \geq t_2)$) and strict inequalities $(t_1 < t_2)$ are rewritten as negated non-strict inequalities $\neg(t_1 \geq t_2)$. [18] Let φ' be the $\mathcal{LA}(\mathbb{Q})$-formula obtained by φ after these substitutions.

In SMT2LGDP$_1$, which is inspired to the polarity-driven CNF conversion of [Plaisted and Greenbaum 1986], we compute the Boolean abstraction φ'^p of φ' (which plays the role of formula ϕ in (4)) and then, for each $\mathcal{LA}$-atom ψ_i occurring positively [resp. negatively] in φ', we add the disjunction $\neg A_i \vee \psi_i$ [resp. $A_i \vee \neg\psi_i$], where A_i is the Boolean atom of φ'^p corresponding to the $\mathcal{LA}$-atom ψ_i.

In SMT2LGDP$_2$, first we compute the CNF-ization of φ' using the MATHSAT5 CNF-izer, and then we encode each non-unit clause $(l_{i1} \vee \ldots \vee l_{in}) \in \varphi'$ as a LGDP disjunction $[Y_i^1 \wedge l_{i1}] \vee \ldots \vee [Y_i^n \wedge l_{in}]$, where $Y_i^1, \ldots, Y_i^n$ are fresh Boolean variables.

Remark 5.2. We decided to provide two different encodings for several reasons. SMT2LGDP$_1$ is a straightforward and very-natural encoding. However, we have verified empirically, and some

[15] http://yices.csl.sri.com/.

[16] GAMS +CPLEX often gives some errors $\leq 10^{-5}$, which we believe are due to the printing floating-point format: e.g., whilst OPTIMATHSAT reports the value `7728125177/2500000000` with infinite-precision arithmetic, GAMS +CPLEX reports it as its floating-point approximation `3.091250e+00`.

[17] http://disi.unitn.it/∼rseba/optimathsat2014.tgz

[18] Here we implicitly assume that the literals $\neg(t_1 = t_2)$, $(t_1 \neq t_2)$ and $(t_1 < t_2)$ occur *positively* in φ; for negative occurrences the encoding is dual.

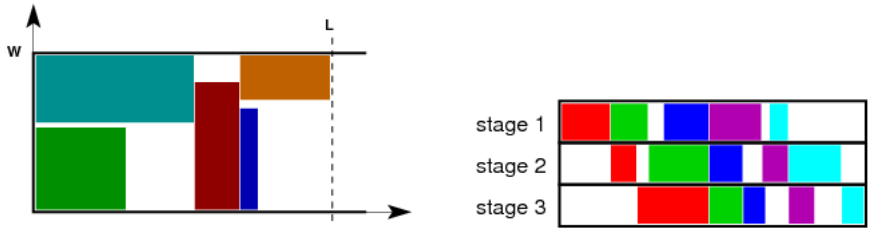

Fig. 2. Graphical representation of a strip-packing (left) and of a zero-wait jobshop problem (right).

discussion with GAMS support team confirmed it [19], that some GAMS tools/options have often problems in handling efficiently and even correctly the Boolean structure of the formulas ϕ in (4) (see e.g. the number of problems terminated with error messages in §5.3-§5.5). Thus, following also the suggestions of the GAMS support team, we have introduced SMT2LGDP$_2$, which eliminates any Boolean structure, reducing the encoding substantially to a set of LGDP disjunctions. Notice, however, that SMT2LGDP$_2$ benefits from the CNF encoder of MATHSAT5.

5.2. Comparison on LGDP problems

We have performed the first comparison over two distinct benchmarks, *strip-packing* and *zero-wait job-shop scheduling* problems, which have been previously proposed as benchmarks for LOGMIP and JAMS by their authors [Vecchietti and Grossmann 2004; Sawaya and Grossmann 2005; Sawaya and Grossmann 2012]. We adopted the encoding of the problems into LGDP given by the authors [20] and gave a corresponding OMT($\mathcal{LA}(\mathbb{Q})$) encoding. We refer to them as "directly generated" benchmarks.

In order to make the results independent from the encoding used, to investigate the correctness and effectiveness of the encodings described in §5.1, and to check the robustness of the tools wrt. different encodings, we also generated formulas from "directly generated" benchmarks by applying the encodings SMT2LGDP$_1$, SMT2LGDP$_2$, and LGDP2SMT; we also applied the SMT2LGDP$_1$/SMT2LGDP$_2$ and LGDP2SMT encodings consecutively to SMT formulas. We refer to them as "encoded" benchmarks.

5.2.1. The strip-packing problem.. Given a set of N rectangles of different length L_i and height H_i, $i \in 1, .., N$, and a strip of fixed width W but unlimited length, the *strip-packing* problem aims at minimizing the length L of the filled part of the strip while filling the strip with all rectangles, without any overlap and any rotation. (See Figure 2 left.)

The LGDP model provided by [Sawaya and Grossmann 2005] is the following:

$$
\begin{array}{lcl}
\min & L & \\
\text{s.t.} & L \geq x_i + L_i & \forall i \in N \\
& \left[\begin{array}{c} Y^1_{ij} \\ x_i + L_i \leq x_j \end{array}\right] \vee \left[\begin{array}{c} Y^2_{ij} \\ x_j + L_j \leq x_i \end{array}\right] & \\
& \vee \left[\begin{array}{c} Y^3_{ij} \\ y_i - H_i \geq y_j \end{array}\right] \vee \left[\begin{array}{c} Y^4_{ij} \\ y_j - H_j \geq y_i \end{array}\right] & \forall i, j \in N, i < j \\
& x_i \leq \mathsf{ub} - L_i & \forall i \in N \\
& H_i \leq y_i \leq W & \forall i \in N \\
& L, x_i, y_i \in \mathbb{R}^1_+, Y^1_{ij}, Y^2_{ij}, Y^3_{ij}, Y^4_{ij} \in \{True, False\} &
\end{array}
\tag{14}
$$

[19] GAMS support team, email personal communication, 2012.
[20] Examples are available at http://www.logmip.ceride.gov.ar/newer.html and at http://www.gams.com/modlib/modlib.htm.

Procedure	Strip-packing													
	$W = \sqrt{N}/2$						$W = 1$						Total	
	$N = 9$		$N = 12$		$N = 15$		$N = 9$		$N = 12$		$N = 15$			
	#s.	time	#s.	time	#s.	time	#s.	time	#s.	time	#s.	time	#s.	time
Directly Generated Benchmarks														
OM-LIN-OF	100	53	100	605	94	8160	100	749	89	3869	54	5547	537	18983
OM-LIN-IN	100	12	100	144	100	3518	100	173	94	2127	74	6808	**568**	**12782**
OM-BIN-OF	100	50	100	625	89	8346	100	588	89	5253	45	5611	523	20473
OM-BIN-IN	100	14	100	211	98	4880	100	202	94	2985	65	8101	557	16393
OM-ADA-IN	100	13	100	192	99	5574	100	214	94	2675	63	7949	556	16617
JAMS(BM)	100	230	78	10177	12	1180	100	158	91	3878	51	6695	**432**	**22318**
JAMS(CH)	100	2854	27	2393	1	417	100	1906	70	7471	17	4032	315	19073
LogMIP(BM)	100	229	78	10159	12	1192	100	157	91	3866	51	6720	**432**	**22323**
LogMIP(CH)	100	2851	27	2414	1	424	100	1907	70	7440	17	4037	315	19073
LGDP2SMT Encoded Benchmarks														
OM-LIN-IN	100	12	100	144	100	3563	100	183	94	2169	73	6466	567	12537
SMT2LGDP$_1$-LGDP2SMT Encoded Benchmarks														
OM-LIN-IN	100	13	100	166	100	5919	100	195	94	2156	74	7080	568	15529
SMT2LGDP$_2$-LGDP2SMT Encoded Benchmarks														
OM-LIN-IN	100	13	100	141	100	5574	100	172	94	2148	74	6650	568	12618
SMT2LGDP$_1$ Encoded Benchmarks														
JAMS(BM)	100	389	68	8733	12	1934	100	162	89	5565	47	7313	**416**	**24096**
JAMS(CH)	99	980	46	6099	2	769	100	726	72	7454	17	3505	336	19533
LogMIP(BM)	100	390	68	8723	12	1946	100	163	89	5547	47	7299	**416**	**24068**
LogMIP(CH)	99	981	54	5480	12	735	100	725	74	7433	17	3542	346	18896
SMT2LGDP$_2$ Encoded Benchmarks														
JAMS(BM)	100	190	81	8460	11	2066	100	159	89	2960	56	8142	437	21977
JAMS(CH)	98	3799	24	2137	1	292	100	2402	68	7926	16	3429	307	19985
LogMIP(BM)	100	191	81	8462	11	2071	100	159	90	2964	56	8206	**438**	**22053**
LogMIP(CH)	98	3807	24	2133	1	312	100	2388	68	7915	17	4027	308	20582

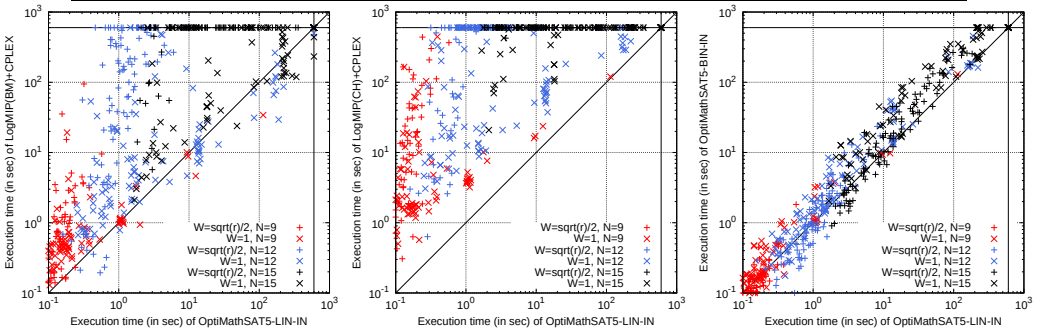

Fig. 3. Table: results (# of solved instances, cumulative time in seconds for solved instances) for OptiMathSAT and GAMS (using LogMIP and JAMS) on 100 random instances (including "directly generated" and "encoded" benchmarks) each of the strip-packing problem for N rectangles, where $N = 9, 12, 15$, and width $W = \sqrt{N}/2, 1$. (We use "OM" as shortcut for OptiMathSAT and omit "+Cplex" in the labels of GAMS tools.) Values highlighted in **bold** represent best performances. Scatter-plots: comparison of the best configuration of OptiMathSAT (OptiMathSAT-LIN-IN) against LogMIP(BM)+CPLEX (left), LogMIP(CH)+CPLEX (center) and OptiMathSAT-BIN-IN (right) on "directly generated" benchmarks.

where L corresponds to the objective function to minimize and every rectangle $j \in J$ is represented by the constants L_j and H_j (length and height respectively) and the variables x_j, y_j (the coordinates of the upper left corner in the 2-dimensional space). Every pair of rectangles $i, j \in N, i < j$ is constrained by a disjunction that avoids their overlapping (each disjunct represents the position of rectangle i in relation to rectangle j). The size of the strip limits the position of each rectangle j: the width of the strip W and the upper bound ub on the optimal solution bound the y_j-coordinate and the height H_j bounds the x_j-coordinate. We express straightforwardly the LGDP model (14) into

OMT($\mathcal{LA}(\mathbb{Q})$) as follows:

$$\begin{array}{rcl}\varphi & \stackrel{\text{def}}{=} & (\mathsf{cost} = L) \wedge \bigwedge_{i\in N}(L \geq x_i + L_i) \\ & \wedge & \bigwedge_{i,j\in N, i<j}\Big((x_i + L_i \leq x_j) \vee (x_j + L_j \leq x_i) \\ & & \vee (y_i - H_i \geq y_j) \vee (y_j - H_j \geq y_i)\Big) \\ & \wedge & \bigwedge_{i\in N}(x_i \leq \mathsf{ub} - L_i) \wedge \bigwedge_{i\in N}(x_i \geq 0) \\ & \wedge & \bigwedge_{i\in N}(H_i \leq y_i) \wedge \bigwedge_{i\in N}(W \geq y_i) \wedge \bigwedge_{i\in N}(y_i \geq 0)\end{array} \tag{15}$$

We randomly generated instances of the strip-packing problem according to a fixed width W of the strip and a fixed number of rectangles N. For each rectangle $j \in N$, length L_j and height H_j are selected in the interval $]0, 1]$ uniformly at random. The upper bound ub is computed with the same heuristic used by [Sawaya and Grossmann 2005], which sorts the rectangles in non-increasing order of width and fills the strip by placing each rectangles in the bottom-left corner, and the lower bound lb is set to zero. We generated 100 samples each for 9, 10 and 11 rectangles and for two values of the width $\sqrt{N}/2$ and 1 (Notice that with $W = \sqrt{N}/2$ the filled strip looks approximatively like a square, whilst $W = 1$ is half the average size of one rectangle.)

5.2.2. The zero-wait jobshop problem.. Consider the scenario where there is a set I of jobs which must be scheduled sequentially on a set J of consecutive stages with zero-wait transfer between them. Each job $i \in I$ has a start time s_i and a processing time t_{ij} in the stage $j \in J_i$, J_i being the set of stages of job i. The goal of the *zero-wait job-shop scheduling* problem is to minimize the makespan, that is the total length of the schedule. (See Figure 2 right.)

The LGDP model provided by [Sawaya and Grossmann 2005] is:

$$\begin{array}{lcl}\min & M & \\ \text{s.t.} & M \geq \sum_{j\in J_i} t_{ij} & \forall i \in I \\ & \Big[\begin{array}{c} Y^1_{ik} \\ s_i + \sum_{m\in J_i, m\leq j} t_{im} \leq s_k + \sum_{m\in J_k, m<j} t_{km}\end{array}\Big] \vee & \\ & \Big[\begin{array}{c} Y^2_{ik} \\ s_k + \sum_{m\in J_k, m\leq j} t_{km} \leq s_i + \sum_{m\in J_i, m<j} t_{im}\end{array}\Big] & \forall j \in C_{ik}, \forall i, k \in I, i < k \\ & M, s_i \in \mathbb{R}^1_+, Y^1_{ik}, Y^2_{ik} \in \{True, False\} & \forall i, k \in I, i < k\end{array} \tag{16}$$

where M corresponds to the objective function to minimize and every job $i \in I$ is represented by the variable s_i (its start time) and the constant t_{ij} (its processing time in stage $j \in J_i$). For each pair of jobs $i, k \in I$ and for each stage j with potential clashes (i.e $j \in C_{ik} = \{J_i \cap J_k\}$), a disjunction ensures that no clash between jobs occur at any stage at the same time. We encoded the corresponding LGDP model (16) into OMT($\mathcal{LA}(\mathbb{Q})$) as follows:

$$\begin{array}{rcl}\varphi & \stackrel{\text{def}}{=} & (\mathsf{cost} = M) \wedge \bigwedge_{i\in I}(M \geq s_i + \sum_{j\in J_i} t_{ij}) \wedge \bigwedge_{i\in I}(s_i \geq 0) \\ & \wedge & \bigwedge_{j\in C_{ik}, \forall i,k\in I, i<k}\Big((s_i + \sum_{m\in J_i, m\leq j} t_{im} \leq s_k + \sum_{m\in J_k, m<j} t_{km}) \\ & & \vee (s_k + \sum_{m\in J_k, m\leq j} t_{km} \leq s_i + \sum_{m\in J_i, m<j} t_{im})\Big)\end{array} \tag{17}$$

We generated randomly instances of the zero-wait jobshop problem according to a fixed number of jobs I and a fixed number of stages J. For each job $i \in I$, start time s_i and processing time t_{ij} of every job are selected in the interval $]0, 1]$ uniformly at random. We consider a set of 100 samples each for 9, 10, 11, 12 jobs and 8 stages and for 11 jobs and 9, 10 stages. We set no value for ub and $\mathsf{lb} = 0$.

5.2.3. Discussion. The table of Figure 3 shows the number of solved instances and their cumulative execution time for different configurations of OPTIMATHSAT and GAMS on "directly generated" and "encoded" benchmarks. The scatter-plots of Figure 3 compare the best-performing

Procedure	Job-shop													
	$I=9$, $J=8$		$I=10$, $J=8$		$I=11$, $J=8$		$I=12$, $J=8$		$I=11$, $J=9$		$I=11$, $J=10$		Total	
	#s.	time	#s.	time	#s.	time	#s.	time	#s.	time	#s.	time	#s.	time
Directly Generated Benchmarks														
OM5-LIN-OF	100	386	100	1854	97	9396	57	14051	100	9637	99	10670	553	45995
OM5-LIN-IN	100	317	100	1584	100	8100	77	18046	100	7738	100	7433	**577**	**43228**
OM5-BIN-OF	100	726	100	3817	88	13222	38	12529	92	14183	90	13287	508	57764
OM5-BIN-IN	100	602	100	3270	97	12878	54	16234	96	13159	96	12350	543	58493
OM5-ADA-IN	100	596	100	3230	97	12262	53	14810	96	12805	96	12125	542	55828
JAMS(BM)	100	268	100	1113	100	4734	87	17067	100	4941	100	6122	**587**	**34245**
JAMS(CH)	84	23830	4	1596	0	0	0	0	0	0	1	363	89	25789
LogMIP(BM)	100	267	100	1114	100	4718	87	17108	100	4962	100	6174	**587**	**34343**
LogMIP(CH)	84	23871	4	1622	0	0	0	0	0	0	1	338	89	25831
LGDP2SMT Encoded Benchmarks														
OM5-LIN-IN	100	324	100	1571	100	7739	74	16494	100	7175	100	7504	574	40807
SMT2LGDP$_1$-LGDP2SMT Encoded Benchmarks														
OM5-LIN-IN	100	336	100	1578	100	7762	71	16589	100	7726	100	7706	571	41697
SMT2LGDP$_2$-LGDP2SMT Encoded Benchmarks														
OM5-LIN-IN	100	320	100	1533	100	7623	68	15120	100	7216	100	7598	568	39410
SMT2LGDP$_1$ Encoded Benchmarks														
JAMS(BM)	100	239	100	1128	100	5516	84	19949	100	6667	100	4176	**584**	**37675**
JAMS(CH)	100	14527	46	17887	0	0	0	0	1	497	0	0	147	32911
LogMIP(BM)	100	240	100	1122	100	5510	83	19489	100	6684	100	4180	583	37225
LogMIP(CH)	100	14465	47	18206	0	0	0	0	1	495	0	0	148	33166
SMT2LGDP$_2$ Encoded Benchmarks														
JAMS(BM)	100	319	100	1865	100	12470	45	15704	97	13189	96	15773	**538**	**59320**
JAMS(CH)	95	22435	18	8030	2	671	0	0	1	526	3	1043	119	32723
LogMIP(BM)	100	319	100	1871	100	12440	45	15747	98	13661	95	15102	**538**	**59140**
LogMIP(CH)	95	22401	18	7991	1	163	0	0	1	437	3	1020	118	32012

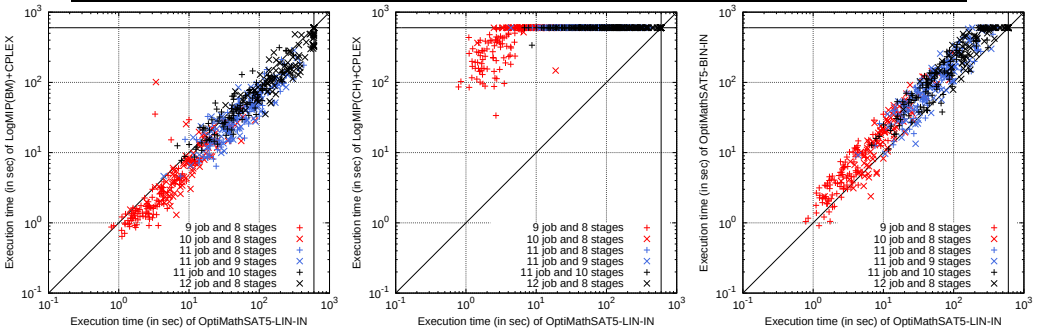
Execution time (in sec) of LogMIP(BM)+CPLEX
Execution time (in sec) of LogMIP(CH)+CPLEX
Execution time (in sec) of OptiMathSAT-BIN-IN
Execution time (in sec) of OptiMathSAT5-LIN-IN
9 job and 8 stages
10 job and 8 stages
11 job and 8 stages
11 job and 9 stages
11 job and 10 stages
12 job and 8 stages

Fig. 4. Table: results (# of solved instances, cumulative time in seconds for solved instances) for OPTIMATHSAT and GAMS on 100 random samples (including "directly generated" and "encoded" benchmarks) each of the job-shop problem for $I = 9, 10, 11, 12$ jobs and $J = 8$ stages and for $I = 11$ jobs and $J = 9, 10$ stage. (We use "OM" as shortcut for OptiMathSAT and omit "+CPLEX" in the labels of GAMS tools.) Values highlighted in **bold** represent best performances. Scatter-plots: comparison of the best configuration of OPTIMATHSAT (OPTIMATHSAT-LIN-IN) against LOGMIP(BM)+CPLEX (left), LOGMIP(CH)+CPLEX (center) and OPTIMATHSAT-BIN-IN (right) on "directly generated" benchmarks.

version of OPTIMATHSAT, OPTIMATHSAT-LIN-IN, against LOGMIP+CPLEX with BM and CH reformulation (left and center respectively) and the two inline versions OPTIMATHSAT-LIN-IN and OPTIMATHSAT-BIN-IN (right) on "directly generated" benchmarks.

The table of Figure 4 shows the number of solved instances and their cumulative execution time for different configurations of OPTIMATHSAT and GAMS on "directly generated" and "encoded" benchmarks. The scatter-plots of Figure 4 compare, on "directly encoded" benchmarks, the best-performing version of OPTIMATHSAT, OPTIMATHSAT-LIN-IN, against LOGMIP with BM and CH reformulation (left and center respectively); the figure also compares the two inline versions OPTIMATHSAT-LIN-IN and OPTIMATHSAT-BIN-IN (right).

The results on the LGDP problems in Figures 3, 4 suggest some considerations.
Comparing the different versions of OPTIMATHSAT, we notice that:

— the inline versions (-IN) behave pairwise uniformly better than the corresponding offline versions (-OF), which is not surprising;
— overall the -LIN options seems to perform a little better than the corresponding -BIN and -ADA options (although gaps are not dramatic).

Remark 5.3. We notice that with LGDP problems binary search is not "obviously faster" than linear search, in compliance with what stated in point 6. in §4.1. This is further enforced by the fact that in strip-packing (15) [resp. job-shop (17)] encodings, the cost variables $\mathsf{cost} \stackrel{\text{def}}{=} L$ [resp. $\mathsf{cost} \stackrel{\text{def}}{=} M$] occurs only in positive unit clauses in the form $(L \geq \langle term \rangle)$ [resp. $(M \geq \langle term \rangle)$]; thus, learning $\neg(\mathsf{cost} < \mathsf{pivot})$ as a result of the binary-search steps with UNSAT results produces no constraining effect on the variables in $\langle term \rangle$, and hence no substantial extra search-pruning effect due to the early-pruning technique of the SMT solver.

Comparing the different versions of the GAMS tools, we see that LOGMIP and JAMS reformulations lead to substantially identical performance on both strip-packing and job-shop instances. For both reformulation tools, the BM versions uniformly outperform the CH ones, often dramatically.
Comparing the performances of the versions of OPTIMATHSAT against these of the GAMS tools, we notice that

— on *strip-packing problems* all versions of OPTIMATHSAT outperform all GAMS versions, regardless of the encoding used. E.g., the best OPTIMATHSAT version solved $\approx 30\%$ more formulas than the best GAMS version;
— on *job-shop problems* results are mixed. OPTIMATHSAT drastically outperforms the CH versions on all encodings and it slightly beats the BM ones on "SMT2LGDP$_2$ encoded" benchmarks, whilst it is slightly beaten by the BM versions on "directly generated" and "SMT2LGDP$_1$ encoded" benchmarks. E.g., the best OPTIMATHSAT version solved $\approx 2\%$ less formulas than the best GAMS version.

Overall, we can conclude that OPTIMATHSAT performances on these problems are comparable with, and most often significantly better than, those of GAMS tools.

We may wonder how these results are affected by the different encodings used. (We recall from the beginning of §5 that all solvers agreed on the results, regardless of the encoding.) In terms of performances, comparing the effects of the different encodings, we notice the following facts.

— On OPTIMATHSAT (-LIN-IN) the effects of the different encodings is substantially negligible, on both strip-packing and job-shop problems, since we have only very small variations in the number of solved instances between "directly generated" and "encoded" instances, in the various encoding combinations. From this reason, we conclude that OPTIMATHSAT is robust wrt. the encodings of these problems.
— On GAMS tools the effects of the different encodings are more relevant, although very heterogeneous: e.g., wrt. to "directly generated" instances, "SMT2LGDP$_1$ encoded" solved formulas are slightly less with BM options, and up to much more with CH options; "SMT2LGDP$_2$ encoded" solved formulas are slightly more on strip-packing and a little less on job-shop with BM options, slightly less on strip-packing and a much more on job-shop with CH options. For this reason, in next sections we always report the results with both encodings.

5.2.4. Analysis of OPTIMATHSAT *performances..* We want to perform a more fine-grained analysis of the performances of the best version of OPTIMATHSAT, OPTIMATHSAT-LIN-IN. To this extent, we partition the total execution time taken on each problem into three consecutive components:

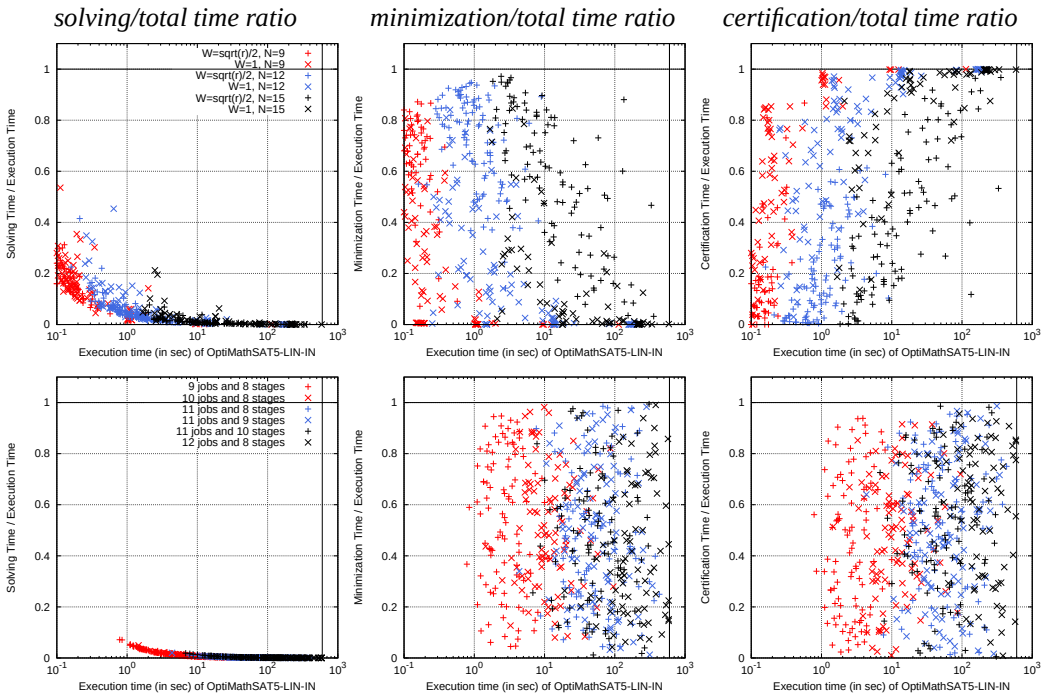

Fig. 5. Scatter-plots comparing solving, minimization and certification time (left, center and right respectively) with the execution time of OPTIMATHSAT-LIN-IN on "directly generated" instances of strip-packing (top) and job-shop (bottom).

— *solving time*, i.e. the time spent on finding the first sub-optimal solution,
— *minimization time*, i.e. the time required to search for the optimal solution,
— and *certification time*, i.e. the time needed for checking there is no better solution.

Figure 5 reports, for all strip-packing (top) and job-shop (bottom) instances, the ratios of the three components above over total execution time. (Notice the log scale of the x axis and the linear scale on the y axis.) We notice a few facts:

— the solving time is nearly negligible, in particular on hardest problems. This tells us, among other facts, that OMT($\mathcal{LA}(\mathbb{Q})$) on these formulas is a much harder problem that plain SMT($\mathcal{LA}(\mathbb{Q})$) on the same formulas;
— the remaining time, on average, is either evenly shared between the minimization and the certification efforts (job-shop, bottom) or even it is mostly dominated by the latter, in particular on the hardest problem (strip-packing, top).

Overall, this suggests that on these instances OPTIMATHSAT-LIN-IN takes on average less than half of the total execution time to find the actual optimal solution, and more than half to prove that there is no better one.

5.3. Comparison on SMT-LIB problems

As a second comparison, in Figure 6 we compare OPTIMATHSAT against the GAMS tools on the satisfiable $\mathcal{LA}(\mathbb{Q})$-formulas (QF_LRA) in the SMT-LIB, augmented with randomly-selected costs. (Hereafter we do not consider the -OF versions of OPTIMATHSAT.) These instances are all classified as "industrial", because they come from the encoding of different real-world problems in formal verification, planning and optimization. They are divided into six categories, namely: `sc`, `uart`, `sal`, `TM`, `tta_startup`, and `miplib`. Notice that other SMT-LIB categories like

Procedure	SMT-LIB/QF_LRA formulas						
	#inst.	#term.	#correct.	#err.msg.	#wrong	#unfeas.	time
OPTIMATHSAT5-LIN-IN	194	194	194	0	0	0	1604
OPTIMATHSAT5-BIN-IN	194	194	**194**	0	0	0	**1449**
OPTIMATHSAT5-ADA-IN	194	194	194	0	0	0	1618
LDGP-SMT-Encoded Benchmarks (SMT2LGDP$_1$ -LGDP2SMT)							
OPTIMATHSAT5-LIN-IN	194	194	194	0	0	0	1820
LDGP-SMT-Encoded Benchmarks (SMT2LGDP$_2$ -LGDP2SMT)							
OPTIMATHSAT5-LIN-IN	194	194	194	0	0	0	1597
LGDP-Encoded Benchmarks (SMT2LGDP$_1$)							
JAMS(BM)+CPLEX	194	171	116	52	0	3	1561
JAMS(CH)+CPLEX	194	193	15	108	0	70	559
LOGMIP(BM)+CPLEX	194	172	**117**	52	0	3	**2152**
LOGMIP(CH)+CPLEX	194	193	15	108	0	70	576
LGDP-Encoded Benchmarks (SMT2LGDP$_2$)							
JAMS(BM)+CPLEX	194	172	**166**	0	4	2	**6839**
JAMS(CH)+CPLEX	194	105	104	0	0	1	9912
LOGMIP(BM)+CPLEX	194	171	165	0	4	2	4103
LOGMIP(CH)+CPLEX	194	105	104	0	0	1	9649

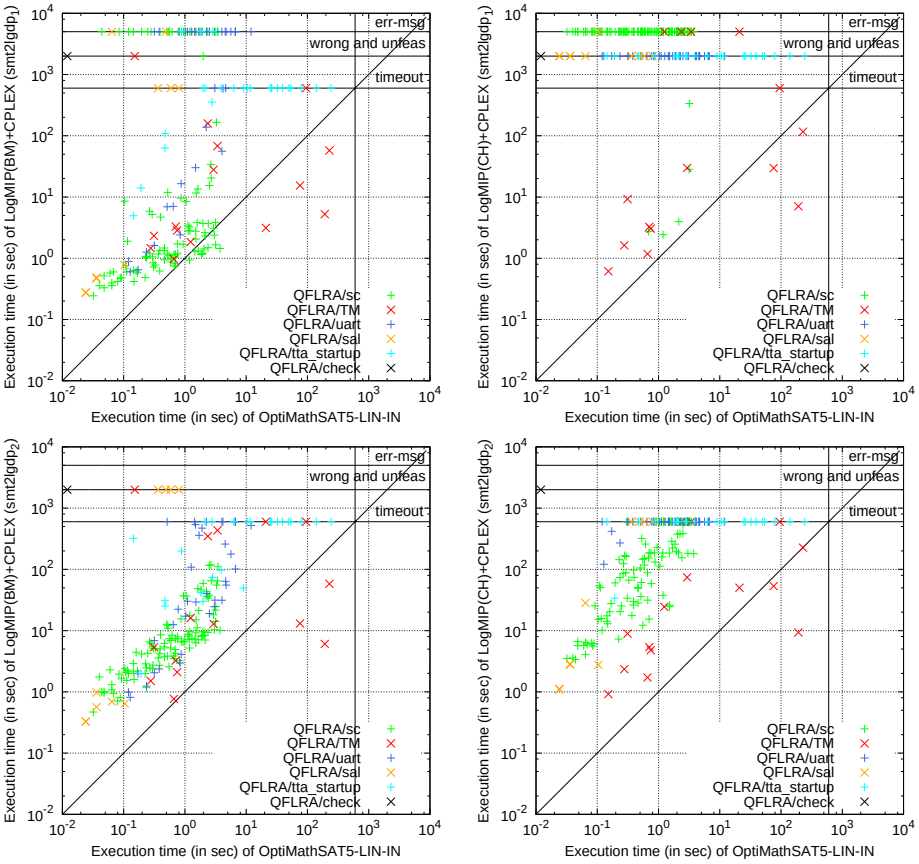

Fig. 6. Table: Results for all the inline versions of OPTIMATHSAT and all the GAMS tools, on a subset of SMT-LIB $\mathcal{LA}(\mathbb{Q})$ satisfiable instances. The columns report respectively: # of instances considered, # of instances terminating within the timeout, # of instances terminating with correct solution, # of instances terminating with error messages, # of instances terminating returning a wrong minimum, # of instances terminating wrongly returning "unfeasible". Scatter-plots: pairwise comparisons on the smt-lib $\mathcal{LA}(\mathbb{Q})$ satisfiable instances between OPTIMATHSAT-LIN-IN and the two versions of LOGMIP+CPLEX. LGDP models are generated using SMT2LGDP$_1$ (top) and SMT2LGDP$_2$ (bottom).

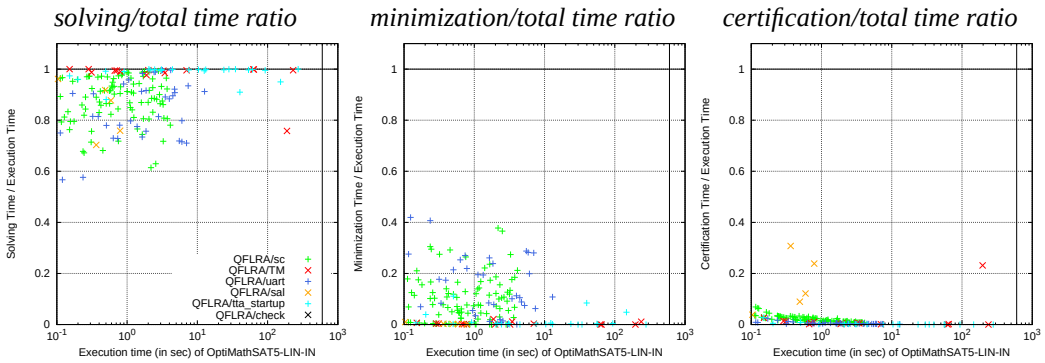

Fig. 7. Scatter-plots comparing solving, minimization and certification time (left, center and right respectively) with the execution time of OPTIMATHSAT-LIN-IN on SMT-LIB instances.

`spider_benchmarks` and `clock_synchro` do not contain satisfiable instances and are thus not reported here.

Since we have no control on the origin of each problem and on the name and meaning of the variables, we selected iteratively one variable at random as cost variable, dropping it if the resulting minimum was $-\infty$. This forced us to eliminate a few instances, in particular all `miplib` ones. We used both SMT2LGDP$_1$ and SMT2LGDP$_2$ to encode these problems into LGDP.

As before, to check for both correctness and effectiveness of the encodings, we also encoded the problems into LGDP by each encoding and encoded then back, to be fed to OPTIMATHSAT-LIN-IN (4th and 5th row). We notice that this caused substantial difference in neither correctness nor efficiency.

We notice first that the results for GAMS tools are affected by correctness problems, with both encodings. Consider the encoding SMT2LGDP$_1$. Out of 194 samples, both GAMS tools with the CH option returned "unfeasible" (i.e. inconsistent) on 70 samples and an error message (regarding some unsatisfied disjunctions) on 108 samples. The two versions with BM returned 3 unfeasible solutions and 52 solutions with error messages. Only 15 samples were solved correctly by GAMS tools with the CH option and 117 (with LOGMIP) or 116 (with JAMS) samples with BM ones, whilst OPTIMATHSAT solved correctly all 194 samples. (We recall that all OPTIMATHSAT results were cross-checked, and that the four GAMS tools were fed with the same files.) With SMT2LGDP$_2$ encoding the number of correctly-solved formulas increases, 104 with CH option and 165 (with LOGMIP) or 166 (with JAMS) with BM; there are no error messages and the number of unfeasible solutions of both GAMS tools with the BM and CH options decreases to 2 and 1 respectively, but the number of solutions with wrong minimum increases to 4 with the BM versions.

Importantly, with both encodings, the results for GAMS tools varied by modifying a couple of parameters from their default value, namely "`eps`" and "`bigM Mvalue`". For example, on the above-mentioned `sal` instance with SMT2LGDP$_1$, with the default values the BM versions returned a wrong minimum value "0", the CH versions returned "unfeasible", whilst OPTIMATHSAT returned the correct minimum value "2"; modifying `eps` and `bigM Mvalue`, the results become unfeasible also with BM options. This highlights the fact that there are indeed some correctness and robustness problems with the GAMS tools, regardless of the encodings used. [21]

5.3.1. Discussion. We conjecture that the problems with the GAMS tools may be caused, at least in part, by the fact that GAMS tools use floating-point rather than infinite-precision arithmetic, and they introduce internally an approximated representation of strict inequalities (see Remark 5.1).

[21] We also isolated a subproblem, small enough to be solved by hand, in which the GAMS tools returned evidently-wrong results, and notified it to the GAMS support team, who reckoned the problem and promised to investigate it eventually.

Notice that, unlike with the LGDP problems in §5.2, SMT-LIB problems do contain occurrences of strict [resp. non-strict] inequalities with positive [resp. negative] polarity.

From the perspective of the efficiency, all versions of OPTIMATHSAT solved correctly all problems within the timeout, the -BIN-IN version performing slightly better than the others; GAMS did not solve many samples (because of timeout, wrong solutions and solutions with error messages). Looking at the scatter-plots, we notice that, with the exception of a few samples, OPTIMATHSAT always outperforms the GAMS tools, often by more than one order magnitude. We notice that on these problems SMT2LGDP$_2$ is generally more effective than SMT2LGDP$_1$ and less prone to errors.

Finally, Figure 7 reports for all SMT-LIB instances the ratios of the solution, minimization and certification times over the total execution time for OPTIMATHSAT-LIN-IN. Unlike with Figure 5, we notice that here the solution time is dominating, the minimization time is significant, and the certification time is nearly negligible. This means that on these instances OPTIMATHSAT-LIN-IN takes on average more than half of its execution time to find the first solution, less than half to find the actual optimal solution, and very little time to prove that there is no better one. We conjecture that this is due to the fact that most satisfiable SMT-LIB instances come from the encoding of formal verification steps of bugged systems which, unlike with the LGDP problems of §5.2, have a limited number of solutions.

5.4. Comparison on SAL problems

As a third comparison, in Figure 8 we compare OPTIMATHSAT against the GAMS tools on $\mathcal{LA}(\mathbb{Q})$-formulas obtained by using the SAL Model Checker on a set of bounded verification problems — Bounded Model Checking (BMC) of invariants [Biere et al. 1999] and K-Induction (K-IND) [Sheeran et al. 2000] — of a well-known parametric timed system, Fisher's Protocol [22].

BMC [resp. K-IND] takes a Finite-State Machine M, an invariant property Ψ and an integer bound k, and produces a propositional formula φ which is satisfiable [resp. unsatisfiable] if and only if there exists a k-step execution violating Ψ [resp. a k-step induction proof that Ψ is always verified]. The approach leverages to real-time systems by producing SMT($\mathcal{LA}(\mathbb{Q})$) formulas rather than purely-propositional ones (see, e.g., [Audemard et al. 2002]).

Fisher's Protocol ensures mutual exclusion among N processes using real-time clocks and a shared variable. The problem is parametric into two positive real values, δ_1 and δ_2, describing the delays of some actions. It is known that mutual exclusion, and other properties included in the SAL model, are verified if and only if $\delta_1 < \delta_2$.

We have produced our OMT($\mathcal{LA}(\mathbb{Q})$) problems as follows. We fixed the value of δ_2 (we chose $\delta_2 = 4$), and then we generated six groups of formulas according to the problem solved (BMC or K-IND) and the property addressed (called `mutex`, `mutual-exclusion`, `time-aux3` and `logical-aux1`). For each group, for increasing values of $N \geq 2$ and for a set of sufficiently-big values of $k \geq k^*$, [23] we used SAL to produce the corresponding parametric SMT($\mathcal{LA}(\mathbb{Q})$) formulas, and asked the tool under test to find the minimum value of δ_1 which made the resulting formula $\mathcal{LA}(\mathbb{Q})$-satisfiable (we knew in advance from the problem that, for k big enough, this value is $\delta_1 = \delta_2 = 4.0$). As before, we used both SMT2LGDP$_1$ and SMT2LGDP$_2$ to encode the OMT($\mathcal{LA}(\mathbb{Q})$) benchmarks into LGDP.

5.4.1. Discussion. The results are presented in Figure 8. The three versions of OPTIMATHSAT solved correctly 385, 382 and 381 out of the 392 samples respectively, OPTIMATHSAT-LIN-IN being the best performer.

Considering the GAMS tools with the encoding SMT2LGDP$_1$, the two tools using BM solved on time and correctly only 4 samples over 392 and returned 19 solutions with error messages and 1

[22] Problems available at http://sal.csl.sri.com/examples.shtml

[23] For BMC, k^* is set to the smallest value of k which makes the formula satisfiable, imposing no upper bound on δ_1; for K-IND, k^* is set to the smallest value of k which makes the formula encoding the inductive step unsatisfiable, imposing $\delta_2 > \delta_1$). In these experiments, k^* ranges from 5 to 10, depending on the problem; also, for each problem, k^* does not depend on N.

Procedure	SAL formulas						
	#inst.	#term.	#correct	#err. msg.	#wrong	#unfeas.	time
OPTIMATHSAT5-LIN-IN	392	385	**385**	0	0	0	**44129**
OPTIMATHSAT5-BIN-IN	392	382	382	0	0	0	45869
OPTIMATHSAT5-ADA-IN	392	381	381	0	0	0	44932
LGDP-Encoded Benchmarks (SMT2LGDP$_1$)							
JAMS(BM)+CPLEX	392	24	**4**	19	1	0	**1096**
JAMS(CH)+CPLEX	392	46	0	0	0	46	0
LOGMIP(BM)+CPLEX	392	24	**4**	19	1	0	**1092**
LOGMIP(CH)+CPLEX	392	46	0	0	0	46	0
LGDP-Encoded Benchmarks (SMT2LGDP$_2$)							
JAMS(BM)+CPLEX	392	28	**14**	0	14	0	**1456**
JAMS(CH)+CPLEX	392	31	2	0	0	29	122
LOGMIP(BM)+CPLEX	392	28	**14**	0	14	0	**1428**
LOGMIP(CH)+CPLEX	392	31	2	0	0	29	120

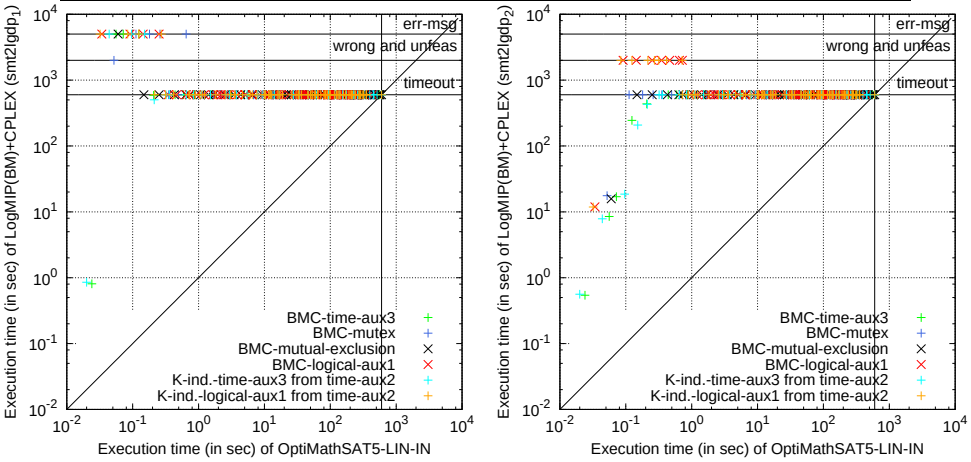

Fig. 8. Table: Results for all the inline versions of OPTIMATHSAT and all the GAMS tools, on formulas generated from SAL models of Fisher's protocol. The columns report respectively: # of instances considered, # of instances terminating within the timeout, # of instances terminating with correct solution, # of instances terminating with error messages (GAMS tools only), # of instances terminating returning a wrong minimum, # of instances terminating wrongly returning "unfeasible". Scatter-plots: comparison of the best configuration of OPTIMATHSAT (OPTIMATHSAT-LIN-IN) against LOGMIP(BM)+CPLEX on SMT2LGDP$_1$ and SMT2LGDP$_2$ encodings (left and right respectively).

solution with wrong minimum, whilst the CH ones always returned "unfeasible". (We recall that all GAMS tools and options are fed the same inputs.) Considering the encoding SMT2LGDP$_2$, the GAMS tools solved more problems correctly (14 with BM tools and 2 with CH), but they returned wrong and unfeasible solutions (14 wrong solutions for BM versions and 29 unfeasible for CH ones). No solution with error messages was found.

The scatter-plots compare OPTIMATHSAT-LIN-IN with the best versions of GAMS, LOGMIP(BM)+CPLEX, on both the encodings, showing that the former dramatically outperforms the latter, no matter the encoding used.

5.5. Comparison on pseudo-Boolean SMT problems

As a fourth comparison, in Figure 9 we evaluate OPTIMATHSAT on the problem sets used in [Cimatti et al. 2010] against the usual GAMS tools and against a recent reimplementation on MATHSAT5 of the tool in [Cimatti et al. 2010], namely PB-MATHSAT, for SMT with Pseudo-

Procedure	MaxSMT / SMT+PB problems						
	#inst.	#term.	#correct.	#err.msg.	#wrong	#unfeas.	time
PB-MATHSAT-LIN	675	636	**636**	0	0	0	**19675**
PB-MATHSAT-BIN	675	632	632	0	0	0	13024
OMT($\mathcal{LA}(\mathbb{Q})$)-Encoded Benchmarks							
OPTIMATHSAT-LIN-IN	675	630	630	0	0	0	20744
OPTIMATHSAT-BIN-IN	675	634	634	0	0	0	16502
OPTIMATHSAT-ADA-IN	675	637	**637**	0	0	0	**18588**
LGDP-Encoded Benchmarks (SMT2LGDP$_1$)							
JAMS(BM)+CPLEX	675	509	19	423	68	8	420
JAMS(CH)+CPLEX	675	642	0	233	41	377	0
LOGMIP(BM)+CPLEX	675	510	**19**	424	68	8	**403**
LOGMIP(CH)+CPLEX	675	642	0	233	41	377	0
LGDP-Encoded Benchmarks (SMT2LGDP$_2$)							
JAMS(BM)+CPLEX	675	449	**92**	9	351	6	**1575**
JAMS(CH)+CPLEX	675	386	48	9	336	2	644
LOGMIP(BM)+CPLEX	675	449	92	9	351	6	1650
LOGMIP(CH)+CPLEX	675	383	48	9	333	2	674

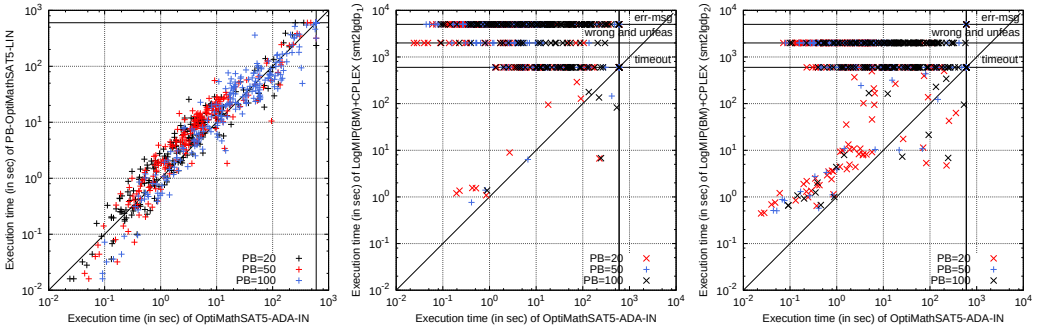

Fig. 9. Table: Results for OPTIMATHSAT, PB-MATHSAT and the GAMS tools, on the MaxSMT benchmarks from [Cimatti et al. 2010]. The columns report respectively: # of instances considered, # of instances terminating within the time-out, # of instances terminating with correct solution, # of instances terminating with error messages (GAMS tools only), # of instances terminating returning a wrong minimum, # of instances terminating wrongly returning "unfeasible". Scatter-plots: comparison of the best configuration of OPTIMATHSAT, OPTIMATHSAT-ADA-IN, against the best configuration of PB-MATHSAT, PB-MATHSAT-LIN (left) and the best configuration of GAMS tools, LOGMIP(BM)+CPLEX, on SMT2LGDP$_1$ and SMT2LGDP$_2$ encodings (center and right respectively).

Boolean constraints (see §3). [24] PB-MATHSAT is tested with both linear search and binary search strategies (denoted with "-LIN" and "-BIN" respectively).

As described in [Cimatti et al. 2010], the problems consists of partial weighted MaxSMT problems which are generated randomly starting from satisfiable $\mathcal{LA}(\mathbb{Q})$-formulas (QF_LRA) in the SMT-LIB, then converted into SMT problems with PB constraints, see (12) in §3. These problems are further encoded into OMT($\mathcal{LA}(\mathbb{Q})$) problems by means of the encoding (11) in §3, and hence into LGDP problems by means of the usual two encodings.

5.5.1. Discussion. The results are presented in Figure 9. The three versions of OPTIMATHSAT solved respectively 630, 634 and 637 problems out of 675 problems overall, whilst the two versions PB-MATHSAT solved respectively 636 and 632. Thus, despite they are both implemented on top of the same SMT solver and PB-MATHSAT is specialized for PB constraints, OPTIMATHSAT performances are analogous to these of the more-specialized tool. The various version of the GAMS

[24] A comparison against the tool in [Cimatti et al. 2010] would not be fair, since the latter was based on the older and slower MATHSAT4. To witness this fact, a comparison of these two implementations is in [Cimatti et al. 2013a].

tool perform drastically worse: with SMT2LGDP$_1$ they solve correctly only a very small number of samples (19 with BM tools and even 0 with CH), returning error messages, unfeasible results or wrong minimum solutions on the remaining set of benchmarks; with SMT2LGDP$_2$ more samples are solved correctly and no error message is produced, but most problems produce a wrong minimum solution.

Remark 5.4. Notice that, unlike with LGDP problems (see Remark 5.3) and in part also with SMT-LIB and SAL problems, with Pseudo-Boolean problems the cost variables occurs in positive unit clauses in the form $(\mathsf{cost} = \langle term \rangle)$; thus, learning $\neg(\mathsf{cost} < \mathsf{pivot})$ as a result of the binary-search steps with UNSAT results produces a constraining effect on the variables in $\langle term \rangle$, and hence a pruning effect in the search due to the early-pruning technique of the SMT solver. This might explain in part the fact that, unlike with previous problems, here binary search performs a little better than linear search.

The scatter-plots in Figure 9 compare the best version of OPTIMATHSAT with these of PB-MATHSAT and of the GAMS tools. We see that OPTIMATHSAT-ADA-IN performances are analogous to these of PB-MATHSAT-LIN , and they are drastically superior to these of GAMS tools with both encodings.

As a side note, in [Cimatti et al. 2013a] another empirical evaluation is performed on MaxSMT problems —although generated with a slightly different random method from SMT-LIB benchmarks— where OPTIMATHSAT performs equivalently better than PB-MATHSAT and the novel specialized MaxSMT tool presented there. We refer the reader to [Cimatti et al. 2013a] for details.

5.6. Comparison against GAMS with parallel CPLEX on all problem sets

As it is common practice in the SMT literature, in this paper we deal with *single-core sequential* procedures. In fact, despite a couple of attempts [Wintersteiger et al. 2009; Kalinnik et al. 2010], the parallelization of SMT-solving procedures is still an open research issue. In particular MATHSAT5, and hence OPTIMATHSAT, provide no support for parallelization. Thus, although the issue of efficiently parallelizing OMT is potentially a very interesting research topic, it is definitely not in the intended scope of this paper.

Unlike with MATHSAT5 and OPTIMATHSAT, however, CPLEX provides full support for multiple-core parallel solving. This is an important benefit, since it allows for exploiting the multiple-core CPUs of current PCs, reducing the elapsed time when searching for a solution. This gives GAMS a potential advantage wrt. OPTIMATHSAT which the previous tests could not reveal.

In order to investigate the actual relevance of this potential advantage, we have recently enriched our empirical investigation by running all the tests in §5.2-§5.5 also on another GAMS tool, namely JAMS(BM)+CPLEX-4CORES: this is the best-performing GAMS GAMS tool in the tests, JAMS(BM)+CPLEX, which uses instead the most recent version of CPLEX, v12.6, *in parallel mode on four cores* (options `opportunistic parallelmode`, `4 threads`). Each of the four threads is given a timeout of 600s.[25] Therefore JAMS(BM)+CPLEX-4CORES is given four times the CPU time resources than its competitors.

5.6.1. Discussion. The results are displayed in Figure 10. In the table, for each group of benchmarks in §5.2-§5.5, we compare the performances of the best-performing OPTIMATHSAT tool, OPTIMATHSAT-LIN-IN, of the best-performing GAMS tool, JAMS(BM)+CPLEX, and of its parallel version, JAMS(BM)+CPLEX-4CORES. The results for the former two tools are taken from Figures 3-9. In the last column we report the mean values of the speedup for JAMS(BM)+CPLEX-4CORES wrt. JAMS(BM)+CPLEX over the problems for which both tools terminated within

[25] We have a technical remark: in order to use the same timeout mechanism for all tools, in all previous tests we have used the Linux command `ulimit` to handle the timeout for all OPTIMATHSAT, GAMS and PB-MATHSAT versions. Unfortunately, `ulimit` does not seem to work properly for multi-threaded processes, so that for JAMS(BM)+CPLEX-4CORES we had to use instead the GAMS/CPLEX internal timeout mechanism, which we have assumed to be reliable.

Procedure	#inst.	#term.	#correct	#err. msg.	#wrong	#unfeas.	time	average speedup
Strip-packing LGDP problems (Directly Generated Benchmarks)								
OPTIMATHSAT-LIN-IN	600	568	**568**	0	0	0	**12782**	–
JAMS(BM)+CPLEX	600	432	432	0	0	0	22318	–
JAMS(BM)+CPLEX-4CORES	600	472	472	0	0	0	17918	5.98
Job-shop LGDP problems (Directly Generated Benchmarks)								
OPTIMATHSAT-LIN-IN	600	577	577	0	0	0	43228	–
JAMS(BM)+CPLEX	600	587	587	0	0	0	34245	–
JAMS(BM)+CPLEX-4CORES	600	600	**600**	0	0	0	**10465**	9.54
SMT-LIB problems (LGDP-Encoded Benchmarks (SMT2LGDP$_2$))								
OPTIMATHSAT5-LIN-IN	194	194	**194**	0	0	0	**1604**	–
JAMS(BM)+CPLEX	194	172	166	0	4	2	6839	–
JAMS(BM)+CPLEX-4CORES	194	155	150	0	4	1	2990	2.36
SAL problems (LGDP-Encoded Benchmarks (SMT2LGDP$_2$))								
OPTIMATHSAT5-LIN-IN	392	385	**385**	0	0	0	**44129**	–
JAMS(BM)+CPLEX	392	28	14	0	14	0	1456	–
JAMS(BM)+CPLEX-4CORES	392	32	17	0	15	0	786	3.66
MaxSMT / SMT+PB problems (LGDP-Encoded Benchmarks (SMT2LGDP$_2$))								
OPTIMATHSAT-LIN-IN	675	630	**630**	0	0	0	**20744**	–
JAMS(BM)+CPLEX	675	449	92	9	351	6	1575	–
JAMS(BM)+CPLEX-4CORES	675	479	95	9	367	6	4033	2.35

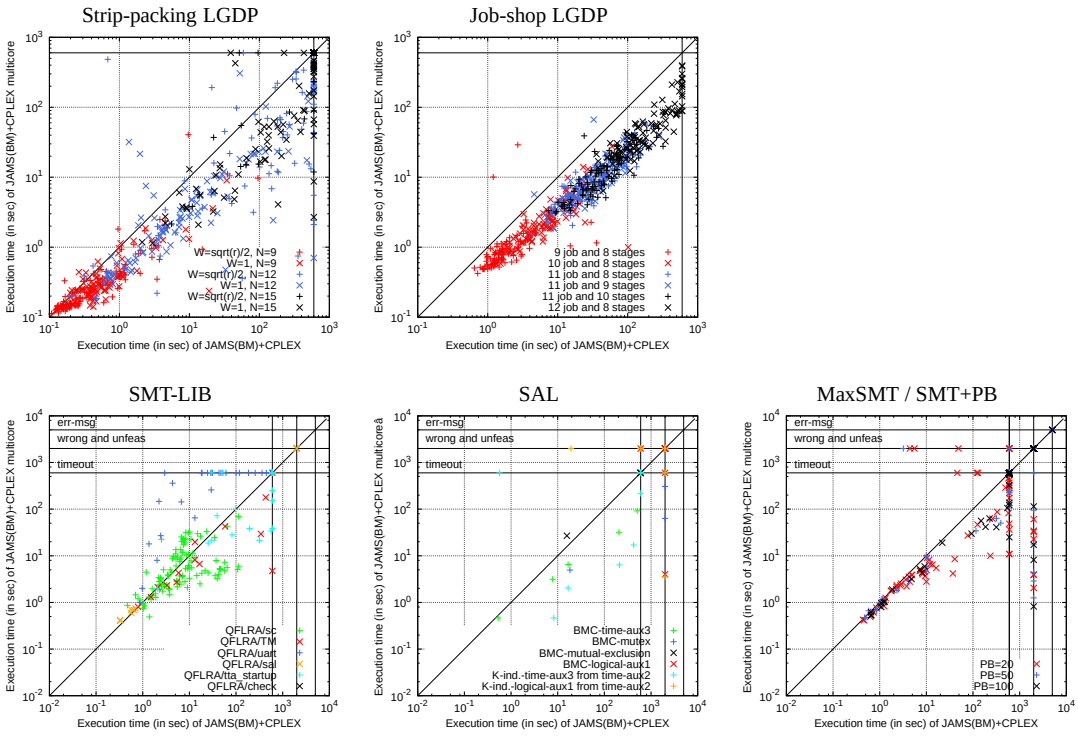

Fig. 10. Table: comparison of OPTIMATHSAT-LIN-IN, JAMS(BM)+CPLEX and JAMS(BM)+CPLEX-4CORES on the five problem sets. Last column: average speedup for JAMS(BM)+CPLEX-4CORES wrt. JAMS(BM)+CPLEX. Scatterplots: pairwise comparison of JAMS(BM)+CPLEX-4CORES vs. JAMS(BM)+CPLEX on the five problem sets.

the timeout. In the scatterplots we compare pairwise the performances of JAMS(BM)+CPLEX-4CORES and JAMS(BM)+CPLEX on the five problem sets.

From Figure 10 we notice the following facts.

— The usage of CPLEX in parallel mode on the four cores pays off in terms of elapsed time: we notice a significant average speedup from JAMS(BM)+CPLEX to JAMS(BM)+CPLEX-4CORES, ranging from 2.35 to 9.54 with the five problem sets.
— The speedup is high and reasonably regular for the two LGDP problem sets, it is lower and quite irregular for the other three problem sets.
— The speedup does not change the qualitative results of the evaluation in the previous sections: OPTIMATHSAT still performs better than all GAMS tools, including JAMS(BM)+CPLEX-4CORES, on the Strip-packing, SMT-LIB, SAL and MaxSMT/SMT+PB problems sets, it performs worse on the Job-shop problem set.

Overall we can conclude that OPTIMATHSAT is very competitive with, and often outperforms, GAMS LGDP tools on the very-extensive set of problems we have used to evaluate them, despite the possibility of GAMS to use CPLEX in parallel mode on multiple-core CPUs. This clearly demonstrates the potential of our novel OMT approach.

6. RELATED WORK.

The idea of optimization in SMT was first introduced by Nieuwenhuis & Oliveras [Nieuwenhuis and Oliveras 2006], who presented a very-general logical framework of "SMT with progressively stronger theories" (e.g., where the theory is progressively strengthened by every new approximation of the minimum cost), and present implementations for MaxSMT based on this framework.

Cimatti et al. [Cimatti et al. 2010] introduced the notion of "Theory of Costs" $\mathcal{C}$ to handle PB cost functions and constraints by an ad-hoc and independent "$\mathcal{C}$-solver" in the standard lazy SMT schema, and implemented a variant of MathSAT tool able to handle SMT with PB constraints and to minimize PB cost functions.

The SMT solvers YICES [Dutertre and Moura 2006] and Z3 [de Moura and Bjørner 2008] also provide support for MaxSMT, although there is no publicly-available document describing the procedures used there.

Ansótegui et al. [Ansótegui et al. 2011] describe the evaluation of an implementation of a MaxSMT procedure based on YICES, although this implementation is not publicly available.

Cimatti et al. [Cimatti et al. 2013a] presented a "modular" approach for MaxSMT, combining a lazy SMT solver with a MaxSAT solver, which can be used as blackboxes.

We recall that MaxSMT and SMT with PB functions can be encoded into each other, and that both are strictly less general than the problem addressed in this paper (§3).

Two other forms of optimization in SMT, which are quite different from the one presented in our work, have been proposed in the literature.

Dillig et al. [Dillig et al. 2012] addressed the problem of finding *partial* models for quantified first-order formulas modulo theories, which minimize the number of free variables which are assigned a value from the domain.Quoting an example from [Dillig et al. 2012], given the formula $\varphi \stackrel{\text{def}}{=} (x + y + w > 0) \vee (x + y + z + w < 5)$, the partial assignment $\{z = 0\}$ satisfies φ because every total assignment extending it satisfies φ and is minimum because there is no assignment satisfying φ which assigns less then one variable. They proposed a general procedure addressing the problem for every theory $\mathcal{T}$ admitting quantifier elimination, and implemented a version for $\mathcal{LA}(\mathbb{Z})$ and $\mathcal{EUF}$ into the MISTRAL tool.

Manolios and Papavasileiou [Manolios and Papavasileiou 2013] proposed the "ILP Modulo Theories" framework as an alternative to SAT Modulo Theories, which allows for combining Integer Linear Programming with decision procedures for signature-disjoint stably-infinite theories $\mathcal{T}$; they presented a general algorithm by integrating the Branch&Cut ILP method with $\mathcal{T}$-specific decision procedures, and implemented it into the INEZ tool. Notice that the approach of [Manolios and Papavasileiou 2013] cannot combine ILP with $\mathcal{LA}(\mathbb{Q})$, since $\mathcal{LA}(\mathbb{Z})$ and $\mathcal{LA}(\mathbb{Q})$ are not signature-disjoint. (See Definition 2 in [Manolios and Papavasileiou 2013].) Also, the objective function is defined on the Integer domain.

We understand that neither of the above-mentioned works can handle the problem addressed in this paper, and vice versa.

7. CONCLUSIONS AND FUTURE WORK

In this paper we have introduced the problem of OMT($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$), an extension of SMT($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$) with minimization of $\mathcal{LA}(\mathbb{Q})$ terms, and proposed two novel procedures addressing it. We have described, implemented and experimentally evaluated this approach, clearly demonstrating all its potentials. We believe that OMT($\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$) and its solving procedures propose as a very-promising tools for a variety of optimization problems.

This research opens the possibility for several interesting future directions. A short-term goal, which we are currently working at, is to extend the approach to $\mathcal{LA}(\mathbb{Z})$ and to mixed $\mathcal{LA}(\mathbb{Q})/\mathcal{LA}(\mathbb{Z})$, by exploiting the solvers which are already present in MATHSAT [Griggio 2012]. As it is implicitly suggested in §5.6, a medium-term goal is to investigate the parallelization of OMT procedures, so that to exploit the power of current multiple-core CPUs. A longer-term goal is to investigate the feasibility of extending the technique to deal with non-linear constraints, possibly using MINLP tools as $\mathcal{T}$-Solver/Minimize.

Acknowledgements

We want to thank the following people for their help and contributions: **Ignacio Grossmann**, **Nicolas Sawaya**, **Aldo Vecchietti**, and the **GAMS support team** for providing help and useful suggestions vie email about LGDP and the usage of the GAMS tools; our colleagues **Alessandro Cimatti**, **Alberto Griggio** and **Patrick Trentin** for their feedback and suggestions; the associate editor and the anonymous reviewers for providing many insightful comments and suggestions.

REFERENCES

ACHTERBERG, T., BERTHOLD, T., KOCH, T., AND WOLTER, K. 2008. Constraint integer programming: a new approach to integrate CP and MIP. In *Proc. CPAIOR'08*. LNCS. Springer, 6–20.

ANSÓTEGUI, C., BOFILL, M., PALAHÍ, M., SUY, J., AND VILLARET, M. 2011. Satisfiability Modulo Theories: An Efficient Approach for the Resource-Constrained Project Scheduling Problem. In *SARA*.

AUDEMARD, G., BOZZANO, M., CIMATTI, A., AND SEBASTIANI, R. 2005. Verifying Industrial Hybrid Systems with MathSAT. In *Proc. BMC 2004*. ENTCS, vol. 119. Elsevier.

AUDEMARD, G., CIMATTI, A., KORNIŁOWICZ, A., AND SEBASTIANI, R. 2002. SAT-Based Bounded Model Checking for Timed Systems. In *Proc. FORTE'02*. LNCS, vol. 2529. Springer.

BALAS, E. 1983. Disjunctive programming and a hierarchy of relaxations for discrete optimization problems. Tech. Rep. DRC-7O-21-R3, Carnegie Mellon University. December. http://repository.cmu.edu/cgi/viewcontent.cgi?article=1949&context=tepper.

BALAS, E. 1998. Disjunctive programming: Properties of the convex hull of feasible points. *Discrete Applied Mathematics 89*, 1-3, 3 – 44.

BALAS, E. AND BONAMI, P. 2007. New variants of lift-and-project cut generation from the lp tableau: Open source implementation and testing. In *IPCO*. 89–103.

BARRETT, C., SEBASTIANI, R., SESHIA, S. A., AND TINELLI, C. 2009. *Satisfiability Modulo Theories*, Chapter 26, 825–885. In Biere et al. Biere et al. [2009].

BARRETT, C. W., DILL, D. L., AND STUMP, A. 2002. A generalization of Shostak's method for combining decision procedures. In *Frontiers of Combining Systems (FROCOS)*. LNAI. Springer-Verlag. Santa Margherita Ligure, Italy.

BIERE, A., CIMATTI, A., CLARKE, E. M., AND ZHU, Y. 1999. Symbolic Model Checking without BDDs. In *Proc. TACAS'99*. 193–207.

BIERE, A., HEULE, M. J. H., VAN MAAREN, H., AND WALSH, T., Eds. 2009. *Handbook of Satisfiability*. Frontiers in Artificial Intelligence and Applications. IOS Press.

BOZZANO, M., BRUTTOMESSO, R., CIMATTI, A., JUNTTILA, T. A., RANISE, S., VAN ROSSUM, P., AND SEBASTIANI, R. 2006. Efficient Theory Combination via Boolean Search. *Information and Computation 204*, 10, 1493–1525.

BROOKE, A., KENDRICK, D., MEERAUS, A., AND RAMAN, R. 2011. *GAMS - A User's Guide*. GAMS Development Corporation, Washington, DC, USA.

CAVADA, R., CIMATTI, A., FRANZÉN, A., KALYANASUNDARAM, K., ROVERI, M., AND SHYAMASUNDAR, R. K. 2007. Computing Predicate Abstractions by Integrating BDDs and SMT Solvers. In *FMCAD*. IEEE Computer Society, 69–76.

CIMATTI, A., FRANZÉN, A., GRIGGIO, A., SEBASTIANI, R., AND STENICO, C. 2010. Satisfiability modulo the theory of costs: Foundations and applications. In *TACAS*. LNCS, vol. 6015. Springer, 99–113.

CIMATTI, A., GRIGGIO, A., SCHAAFSMA, B. J., AND SEBASTIANI, R. 2013a. A Modular Approach to MaxSAT Modulo Theories. In *International Conference on Theory and Applications of Satisfiability Testing, SAT*. LNCS, vol. 7962.

CIMATTI, A., GRIGGIO, A., SCHAAFSMA, B. J., AND SEBASTIANI, R. 2013b. The MathSAT 5 SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems, TACAS'13*. LNCS, vol. 7795. Springer, 95–109.

CIMATTI, A., GRIGGIO, A., AND SEBASTIANI, R. 2010. Efficient Generation of Craig Interpolants in Satisfiability Modulo Theories. *ACM Transaction on Computational Logics 12*, 1 (October).

CIMATTI, A., GRIGGIO, A., AND SEBASTIANI, R. 2011. Computing Small Unsatisfiable Cores in SAT Modulo Theories. *Journal of Artificial Intelligence Research, JAIR 40*, 701–728.

DE MOURA, L. M. AND BJØRNER, N. 2008. Z3: An efficient smt solver. In *TACAS*. LNCS, vol. 4963. Springer, 337–340.

DILLIG, I., DILLIG, T., MCMILLAN, K. L., AND AIKEN, A. 2012. Minimum Satisfying Assignments for SMT. In *CAV*. 394–409.

DUTERTRE, B. AND DE MOURA, L. 2006. A Fast Linear-Arithmetic Solver for DPLL(T). In *CAV*. LNCS, vol. 4144.

DUTERTRE, B. AND MOURA, L. D. 2006. The yices smt solver. Tech. rep., SRI.

EÉN, N. AND SÖRENSSON, N. 2004. An extensible SAT-solver. In *Theory and Applications of Satisfiability Testing (SAT 2003)*. LNCS, vol. 2919. Springer, 502–518.

GRIGGIO, A. 2012. A Practical Approach to Satisfiability Modulo Linear Integer Arithmetic. *Journal on Satisfiability, Boolean Modeling and Computation - JSAT 8*, 1–27.

IBM 2010. *IBM ILOG CPLEX Optimizer*. IBM. Available at http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/.

KALINNIK, N., ÁBRAHÁM, E., SCHUBERT, T., WIMMER, R., AND BECKER, B. 2010. Exploiting Different Strategies for the Parallelization of an SMT Solver. In *Proc. MBMV, Dresden, Germany, February 22-24, 2010*. 97–106.

LI, C. M. AND MANYÀ, F. 2009. *MaxSAT, Hard and Soft Constraints*, Chapter 19, 613–631. In Biere et al. Biere et al. [2009].

LODI, A. 2009. Mixed Integer Programming Computation. In *50 Years of Integer Programming 1958-2008*. Springer-Verlag, 619–645.

LYNCE, I. AND MARQUES-SILVA, J. 2004. On Computing Minimum Unsatisfiable Cores. In *7th International Conference on Theory and Applications of Satisfiability Testing*.

MANOLIOS, P. AND PAPAVASILEIOU, V. 2013. Ilp modulo theories. In *CAV*. 662–677.

MARQUES-SILVA, J. P., LYNCE, I., AND MALIK, S. 2009. *Conflict-Driven Clause Learning SAT Solvers*, Chapter 4, 131–153. In Biere et al. Biere et al. [2009].

MARQUES-SILVA, J. P. AND SAKALLAH, K. A. 1996. GRASP - A new Search Algorithm for Satisfiability. In *Proc. ICCAD'96*. 220–227.

MOSKEWICZ, M. W., MADIGAN, C. F., ZHAO, Y., ZHANG, L., AND MALIK, S. 2001. Chaff: Engineering an Efficient SAT Solver. In *DAC*. 530–535.

NELSON, C. G. AND OPPEN, D. C. 1979. Simplification by cooperating decision procedures. *TOPLAS 1*, 2, 245–257.

NIEUWENHUIS, R. AND OLIVERAS, A. 2006. On SAT Modulo Theories and Optimization Problems. In *Proc. Theory and Applications of Satisfiability Testing - SAT 2006*. LNCS, vol. 4121. Springer.

NIEUWENHUIS, R., OLIVERAS, A., AND TINELLI, C. 2006. Solving SAT and SAT Modulo Theories: from an Abstract Davis-Putnam-Logemann-Loveland Procedure to DPLL(T). *Journal of the ACM 53*, 6 (November), 937–977.

PLAISTED, D. AND GREENBAUM, S. 1986. A Structure-preserving Clause Form Translation. *Journal of Symbolic Computation 2*, 293–304.

RAMAN, R. AND GROSSMANN, I. 1994. Modelling and computational techniques for logic based integer programming. *Computers and Chemical Engineering 18*, 7, 563 – 578.

ROUSSEL, O. AND MANQUINHO, V. 2009. *Pseudo-Boolean and Cardinality Constraints*, Chapter 22, 695–733. In Biere et al. Biere et al. [2009].

SAWAYA, N. W. AND GROSSMANN, I. E. 2005. A cutting plane method for solving linear generalized disjunctive programming problems. *Comput Chem Eng 29*, 9, 1891–1913.

SAWAYA, N. W. AND GROSSMANN, I. E. 2012. A hierarchy of relaxations for linear generalized disjunctive programming. *European Journal of Operational Research 216*, 1, 70–82.

SEBASTIANI, R. 2007. Lazy Satisfiability Modulo Theories. *Journal on Satisfiability, Boolean Modeling and Computation, JSAT 3*, 3-4, 141–224.

SEBASTIANI, R. AND TOMASI, S. 2012. Optimization in SMT with LA(Q) Cost Functions. In *IJCAR*. LNAI, vol. 7364. Springer, 484–498. Available at http://disi.unitn.it/∼rseba/publist.html.

SELLMANN, M. AND KADIOGLU, S. 2008. Dichotomic Search Protocols for Constrained Optimization. In *CP*. LNCS, vol. 5202. Springer.

SHEERAN, M., SINGH, S., AND STÅLMARCK, G. 2000. Checking Safety Properties Using Induction and a SAT-Solver. In *FMCAD*. LNCS. Springer, 108–125.

TINELLI, C. AND HARANDI, M. T. 1996. A new correctness proof of the Nelson–Oppen combination procedure. In *Proc. Frontiers of Combining Systems, FroCoS'06*. Applied Logic. Kluwer.

VECCHIETTI, A. 2011. Personal communication.

VECCHIETTI, A. AND GROSSMANN, I. 2004. Computational experience with logmip solving linear and nonlinear disjunctive programming problems. In *Proc. of FOCAPD*. 587–590.

WINTERSTEIGER, C. M., HAMADI, Y., AND DE MOURA, L. M. 2009. A Concurrent Portfolio Approach to SMT Solving. In *Proc Computer Aided Verification, CAV*. 715–720.

WOLFMAN, S. AND WELD, D. 1999. The LPSAT Engine & its Application to Resource Planning. In *Proc. IJCAI*.

A. APPENDIX: PROOF OF THE THEOREMS

A.1. Proof of Theorem 3.6

We first need proving the following lemmas.

LEMMA A.1. *Let φ be a $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-formula and $\mathcal{E} \stackrel{def}{=} \{\eta_1, \ldots, \eta_n\}$ be the set of all total truth assignments propositionally satisfying φ. Then $\mathsf{min}_{\mathsf{cost}}(\varphi) = min_{\eta_i \in \mathcal{E}} \mathsf{min}_{\mathsf{cost}}(\eta_i)$.*

PROOF. If φ is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-unsatisfiable, then $\mathsf{min}_{\mathsf{cost}}(\varphi) = min_{\eta_i \in \mathcal{E}} \mathsf{min}_{\mathsf{cost}}(\eta_i) = +\infty$. Otherwise, the thesis follows straightforwardly from the fact that the set of the models of φ is the union of the sets of the models of the assignments in $\mathcal{E}$. □

LEMMA A.2. *Let φ be a $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-formula and μ be a $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable* **partial** *assignment s.t. $\mu \models_p \varphi$. Then there exists at least one $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable* **total** *assignment η s.t. $\mu \subseteq \eta$, $\eta \models_p \varphi$, and $\mathsf{min}_{\mathsf{cost}}(\mu) = \mathsf{min}_{\mathsf{cost}}(\eta)$.*

PROOF. Let $\mathcal{I}$ be a model for μ, and hence for φ. Then

$$\eta \stackrel{\text{def}}{=} \bigwedge_{\substack{\psi_i \in Atoms(\varphi) \\ \mathcal{I} \models \psi_i}} \psi_i \wedge \bigwedge_{\substack{\psi_i \in Atoms(\varphi) \\ \mathcal{I} \models \neg\psi_i}} \neg\psi_i \tag{18}$$

By construction, η is a total truth assignment for φ and it is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable, $\mu \subseteq \eta$ and $\mathsf{min}_{\mathsf{cost}}(\eta) = \mathsf{min}_{\mathsf{cost}}(\mu) = \mathcal{I}(\mathsf{cost})$. Since $\mu \subseteq \eta$, then $\eta \models_p \varphi$. □

The proof of Theorem 3.6 then follows.

THEOREM 3.6 1. *Let φ be a $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-formula and let $\mathcal{M} \stackrel{def}{=} \{\mu_1, \ldots, \mu_n\}$ be a complete collection of (possibly partial) truth assignments propositionally satisfying φ. Then $\mathsf{min}_{\mathsf{cost}}(\varphi) = min_{\mu \in \mathcal{M}} \mathsf{min}_{\mathsf{cost}}(\mu)$.*

PROOF. If φ is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-unsatisfiable, then $\mathsf{min}_{\mathsf{cost}}(\varphi) = min_{\mu \in \mathcal{M}} \mathsf{min}_{\mathsf{cost}}(\mu) = +\infty$ by Definition 3.1 and Theorem 3.5. Otherwise, $\mathsf{min}_{\mathsf{cost}}(\varphi) < +\infty$. Then:

Proof of $\mathsf{min}_{\mathsf{cost}}(\varphi) \leq min_{\mu \in \mathcal{M}} \mathsf{min}_{\mathsf{cost}}(\mu)$:
By absurd, suppose exists $\mu \in \mathcal{M}$ s.t. $\mathsf{min}_{\mathsf{cost}}(\mu) < \mathsf{min}_{\mathsf{cost}}(\varphi)$. By Proposition 3.3, μ is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$ satisfiable. By Lemma A.2, there exists a $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable total assignment η s.t. $\mu \subseteq \eta$, $\eta \models_p \varphi$, and $\mathsf{min}_{\mathsf{cost}}(\mu) = \mathsf{min}_{\mathsf{cost}}(\eta)$. By lemma A.1, $\mathsf{min}_{\mathsf{cost}}(\eta) \geq \mathsf{min}_{\mathsf{cost}}(\varphi)$, and hence $\mathsf{min}_{\mathsf{cost}}(\mu) \geq \mathsf{min}_{\mathsf{cost}}(\varphi)$, contradicting the hypothesis.
Proof of $\mathsf{min}_{\mathsf{cost}}(\varphi) \geq min_{\mu \in \mathcal{M}} \mathsf{min}_{\mathsf{cost}}(\mu)$:
From Lemma A.1 we have that $\mathsf{min}_{\mathsf{cost}}(\varphi) = min_{\eta_i \in \mathcal{E}} \mathsf{min}_{\mathsf{cost}}(\eta_i)$. Let $\eta \in \mathcal{E}$ s.t. $\mathsf{min}_{\mathsf{cost}}(\varphi) = \mathsf{min}_{\mathsf{cost}}(\eta) < +\infty$. Hence η is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable. Thus, there exists $\mu \in \mathcal{M}$ s.t. $\mu \subseteq \eta$. μ is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable since η is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable. From Proposition 3.3, $\mathsf{min}_{\mathsf{cost}}(\mu) \leq \mathsf{min}_{\mathsf{cost}}(\eta)$, hence $\mathsf{min}_{\mathsf{cost}}(\mu) \leq \mathsf{min}_{\mathsf{cost}}(\varphi)$. Thus the thesis holds.

□

A.2. Proof of Theorem 3.10

THEOREM 3.10 2. *Let μ be as in Definition 3.9. Then*

(a) $\mathsf{min}_{\mathsf{cost}}(\mu) = min_{\eta \in \mathcal{EX}_{ed}(\mu)} \mathsf{min}_{\mathsf{cost}}(\eta)$

(b) forall $\eta \in \mathcal{EX}_{ed}(\mu)$,

$$\mathsf{min}_{\mathsf{cost}}(\eta) = \begin{cases} +\infty & \text{if } \mu_{\mathcal{T}} \wedge \mu_{ed} \text{ is } \mathcal{T}\text{-unsatisfiable or} \\ & \text{if } \mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ed} \text{ is } \mathcal{LA}(\mathbb{Q})\text{-unsatisfiable} \\ \mathsf{min}_{\mathsf{cost}}(\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ed}) & \text{otherwise.} \end{cases}$$

PROOF.

(*a*) Let

$$\mu' \stackrel{\text{def}}{=} \mu \wedge \bigwedge_{(x_i = x_j) \in \mathcal{IE}(\mu)} ((x_i = x_j) \vee \neg(x_i = x_j)).$$

μ and μ' are obviously $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-equivalent, so that $\mathsf{min}_{\mathsf{cost}}(\mu) = \mathsf{min}_{\mathsf{cost}}(\mu')$. By construction, $\mathcal{EX}_{ed}(\mu)$ is the set of all total truth assignments propositionally satisfying μ', so that $\mathsf{min}_{\mathsf{cost}}(\mu') = min_{\eta \in \mathcal{EX}_{ed}(\mu)} \mathsf{min}_{\mathsf{cost}}(\eta)$.

(*b*) By Theorem 3.8, η is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable if and only if $\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ed}$ is $\mathcal{LA}(\mathbb{Q})$-satisfiable and $\mu_{\mathcal{T}} \wedge \mu_{ed}$ is $\mathcal{T}$-satisfiable. Thus,

- if $\mu_{\mathcal{T}} \wedge \mu_{ed}$ is $\mathcal{T}$-unsatisfiable, then η is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-unsatisfiable, so that $\mathsf{min}_{\mathsf{cost}}(\eta) = +\infty$.
- If $\mu_{\mathcal{T}} \wedge \mu_{ed}$ is $\mathcal{T}$-satisfiable and $\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ed}$ is $\mathcal{LA}(\mathbb{Q})$-unsatisfiable, then η is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-unsatisfiable, so that $\mathsf{min}_{\mathsf{cost}}(\eta) = \mathsf{min}_{\mathsf{cost}}(\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ed}) = +\infty$.
- If $\mu_{\mathcal{T}} \wedge \mu_{ed}$ is $\mathcal{T}$-satisfiable and $\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ed}$ is $\mathcal{LA}(\mathbb{Q})$-satisfiable, then η is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable. We split the proof into two parts.
  - $\leq$ **case:** Let $c \in \mathbb{Q}$ be the value of $\mathsf{min}_{\mathsf{cost}}(\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ed})$. Let $\mu' \stackrel{\text{def}}{=} \mu \wedge (\mathsf{cost} = c)$. Since $(\mathsf{cost} = c)$ is a $\mathcal{LA}(\mathbb{Q})$-pure atom, then $\mu' = \mu'_{\mathcal{T}} \wedge \mu'_{\mathcal{LA}(\mathbb{Q})}$ s.t. $\mu'_{\mathcal{T}} = \mu_{\mathcal{T}}$ and $\mu'_{\mathcal{LA}(\mathbb{Q})} = \mu_{\mathcal{LA}(\mathbb{Q})} \wedge (\mathsf{cost} = c)$, which are respectively $\mathcal{T}$- and $\mathcal{LA}(\mathbb{Q})$-pure and $\mathcal{T}$- and $\mathcal{LA}(\mathbb{Q})$-satisfiable by construction. Let $\eta' \stackrel{\text{def}}{=} \eta \wedge (\mathsf{cost} = c)$. Since $\mathcal{IE}(\mu) = \mathcal{IE}(\mu')$, then μ', $\mu'_{\mathcal{LA}(\mathbb{Q})}$, $\mu'_{\mathcal{T}}$ and η' match the hypothesis of Theorem 3.8, from which we have that η' is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable, so that η has a model $\mathcal{I}$ s.t. $\mathcal{I}(\mathsf{cost}) = c$. Thus, we have that $\mathsf{min}_{\mathsf{cost}}(\eta) \leq \mathsf{min}_{\mathsf{cost}}(\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ed})$.
  - $\geq$ **case:** Let $c \in \mathbb{Q}$ be the value of $\mathsf{min}_{\mathsf{cost}}(\eta)$. Then $\eta \wedge (\mathsf{cost} = c)$ is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable. We define μ', $\mu'_{\mathcal{LA}(\mathbb{Q})}$, $\mu'_{\mathcal{T}}$ and η' as in the "$\leq$" case. As before, they match the hypothesis of Theorem 3.8, from which we have that $\mu'_{\mathcal{LA}(\mathbb{Q})}$ is $\mathcal{LA}(\mathbb{Q})$-satisfiable. Hence, $\mu_{\mathcal{LA}(\mathbb{Q})}$ has a model $\mathcal{I}$ s.t. $\mathcal{I}(\mathsf{cost}) = c$. Thus, we have that $\mathsf{min}_{\mathsf{cost}}(\eta) \geq \mathsf{min}_{\mathsf{cost}}(\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_{ed})$.

□

A.3. Proof of Theorem 3.12

THEOREM 3.12 3. *Let μ be as in Definition 3.9. Then*

(a) *μ is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable iff some $\rho \in \mathcal{EX}_{edi}(\mu)$ is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable.*

(b) *$\mathsf{min}_{\mathsf{cost}}(\mu) = min_{\rho \in \mathcal{EX}_{edi}(\mu)} \mathsf{min}_{\mathsf{cost}}(\rho)$.*

(c) *forall $\rho \in \mathcal{EX}_{edi}(\mu)$, ρ is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable iff $\mu_{\mathcal{T}} \wedge \mu_{ed}$ is $\mathcal{T}$-satisfiable and $\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_e \wedge \mu_i$ is $\mathcal{LA}(\mathbb{Q})$-satisfiable.*

(d) *forall $\rho \in \mathcal{EX}_{edi}(\mu)$,*

$$\mathsf{min}_{\mathsf{cost}}(\rho) = \begin{cases} +\infty & \text{if } \mu_{\mathcal{T}} \wedge \mu_{ed} \text{ is } \mathcal{T}\text{-unsatisfiable or} \\ & \text{if } \mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_e \wedge \mu_i \text{ is } \mathcal{LA}(\mathbb{Q})\text{-unsatisfiable} \\ \mathsf{min}_{\mathsf{cost}}(\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_e \wedge \mu_i) & \text{otherwise.} \end{cases}$$

PROOF. Let

$$\mu^* \stackrel{\text{def}}{=} \mu \wedge \bigwedge_{(x_i = x_j) \in \mathcal{IE}(\mu)} \left(\begin{array}{l} ((x_i = x_j) \vee (x_i < x_j) \vee (x_i > x_j)) \wedge \\ (\neg(x_i = x_j) \vee \neg(x_i < x_j)) \wedge \\ (\neg(x_i = x_j) \vee \neg(x_i > x_j)) \wedge \\ (\neg(x_i < x_j) \vee \neg(x_i > x_j)) \end{array} \right) \tag{19}$$

All clauses in the right conjuncts in (19) are $\mathcal{LA}(\mathbb{Q})$-valid, hence μ and μ^* are $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-equivalent, so that $\mathsf{min}_{\mathsf{cost}}(\mu) = \mathsf{min}_{\mathsf{cost}}(\mu^*)$. By construction, $\mathcal{EX}_{edi}(\mu)$ is the set of all total truth assignments propositionally satisfying μ^*.

(*a*) By Theorem 3.5, μ^* is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable iff some $\rho \in \mathcal{EX}_{edi}(\mu)$ is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable, from which the thesis.

(*b*) $\mathsf{min}_{\mathsf{cost}}(\mu) = \mathsf{min}_{\mathsf{cost}}(\mu^*) = min_{\rho \in \mathcal{EX}_{edi}(\mu)} \mathsf{min}_{\mathsf{cost}}(\rho)$.

(*c*) We consider one $\rho \in \mathcal{EX}_{edi}(\mu)$. $\rho = \mu_{\mathcal{T}} \wedge \mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_e \wedge \mu_d \wedge \mu_i$. We notice that all literals in μ_i are $\mathcal{LA}(\mathbb{Q})$-pure, s.t. it is the $\mathcal{LA}(\mathbb{Q})$-pure part of ρ (namely, $\rho_{\mathcal{LA}(\mathbb{Q})}$). Thus, by Theorem 3.8, ρ is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable iff $\overbrace{\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_i}^{\rho_{\mathcal{LA}(\mathbb{Q})}} \wedge \overbrace{\mu_e \wedge \mu_d}^{\mu_{ed}}$ is $\mathcal{LA}(\mathbb{Q})$-satisfiable and $\mu_{\mathcal{T}} \wedge \overbrace{\mu_e \wedge \mu_d}^{\mu_{ed}}$ is $\mathcal{T}$-satisfiable. By construction, $\mu_i \models_{\mathcal{LA}(\mathbb{Q})} \mu_d$. Thus, $\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_i \wedge \mu_e \wedge \mu_d$ is $\mathcal{LA}(\mathbb{Q})$-satisfiable iff $\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_i \wedge \mu_e$ is $\mathcal{LA}(\mathbb{Q})$-satisfiable. Thus the thesis holds.

(*d*) We consider one $\rho \in \mathcal{EX}_{edi}(\mu)$ and partition it as in point (c). From point (c), if $\mu_{\mathcal{T}} \wedge \mu_{ed}$ is $\mathcal{T}$-unsatisfiable or $\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_i \wedge \mu_e$ is $\mathcal{LA}(\mathbb{Q})$-unsatisfiable, then ρ is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-unsatisfiable, so that $\mathsf{min}_{\mathsf{cost}}(\rho) = +\infty$. Otherwise, ρ is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable.

- $\leq$ **case:** Let $c \in \mathbb{Q}$ be the value of $\mathsf{min}_{\mathsf{cost}}(\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_e \wedge \mu_i)$. Let $\mu' \stackrel{\text{def}}{=} \mu \wedge (\mathsf{cost} = c)$. Since $(\mathsf{cost} = c)$ is a $\mathcal{LA}(\mathbb{Q})$-pure atom, then $\mu' = \mu'_{\mathcal{T}} \wedge \mu'_{\mathcal{LA}(\mathbb{Q})}$ s.t. $\mu'_{\mathcal{T}} = \mu_{\mathcal{T}}$ and $\mu'_{\mathcal{LA}(\mathbb{Q})} = \mu_{\mathcal{LA}(\mathbb{Q})} \wedge (\mathsf{cost} = c)$, which are respectively $\mathcal{T}$- and $\mathcal{LA}(\mathbb{Q})$-pure. Also, $\mu'_{\mathcal{T}} \wedge \mu_{ed}$ is $\mathcal{T}$-satisfiable and $\mu'_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_e \wedge \mu_i$ is $\mathcal{LA}(\mathbb{Q})$-satisfiable by construction. Let $\rho' \stackrel{\text{def}}{=} \rho \wedge (\mathsf{cost} = c)$. Since $\mathcal{IE}(\mu) = \mathcal{IE}(\mu')$, then also μ', $\mu'_{\mathcal{LA}(\mathbb{Q})}$, $\mu'_{\mathcal{T}}$ and ρ' match the hypothesis of this theorem. Thus, by point (c), ρ' is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable, so that ρ has a model $\mathcal{I}$ s.t. $\mathcal{I}(\mathsf{cost}) = c$. Therefore we have that $\mathsf{min}_{\mathsf{cost}}(\rho) \leq \mathsf{min}_{\mathsf{cost}}(\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_e \wedge \mu_i)$.
- $\geq$ **case:** Let $c \in \mathbb{Q}$ be the value of $\mathsf{min}_{\mathsf{cost}}(\rho)$. Then $\rho \wedge (\mathsf{cost} = c)$ is $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{T}$-satisfiable. We define μ', $\mu'_{\mathcal{LA}(\mathbb{Q})}$, $\mu'_{\mathcal{T}}$ and ρ' as in the "$\leq$" case. As before, they also match the hypothesis of this theorem, so that by point (c) $\mu'_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_e \wedge \mu_i$ is $\mathcal{LA}(\mathbb{Q})$-satisfiable. Thus, $\mu'_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_e \wedge \mu_i$ has a model $\mathcal{I}$ s.t. $\mathcal{I}(\mathsf{cost}) = c$. Therefore we have that $\mathsf{min}_{\mathsf{cost}}(\rho) \geq \mathsf{min}_{\mathsf{cost}}(\mu_{\mathcal{LA}(\mathbb{Q})} \wedge \mu_e \wedge \mu_i)$.

□

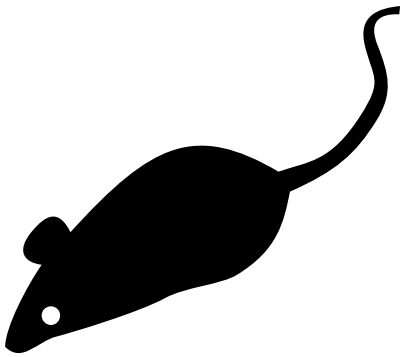

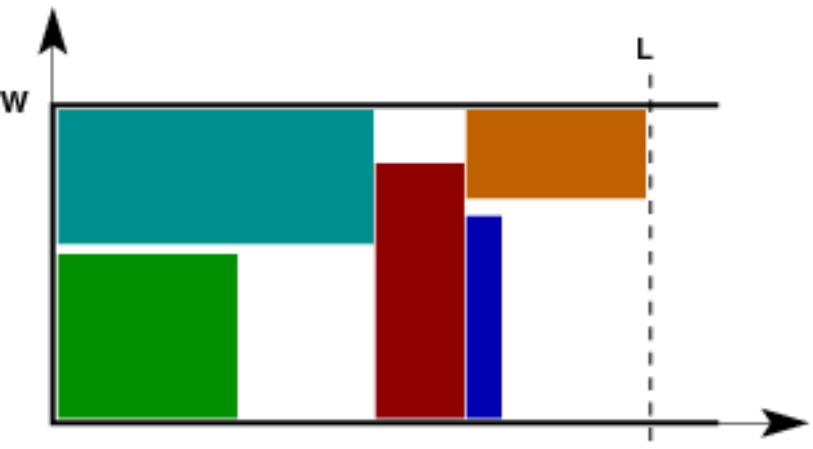
W
L